MIT Technology Review

VOL. 118 NO. 3 MAY/JUNE 2015 $6.99

WE CAN NOW ENGINEER THE HUMAN RACE

p26

INNOVATION

THIS IS WHAT WE DO.

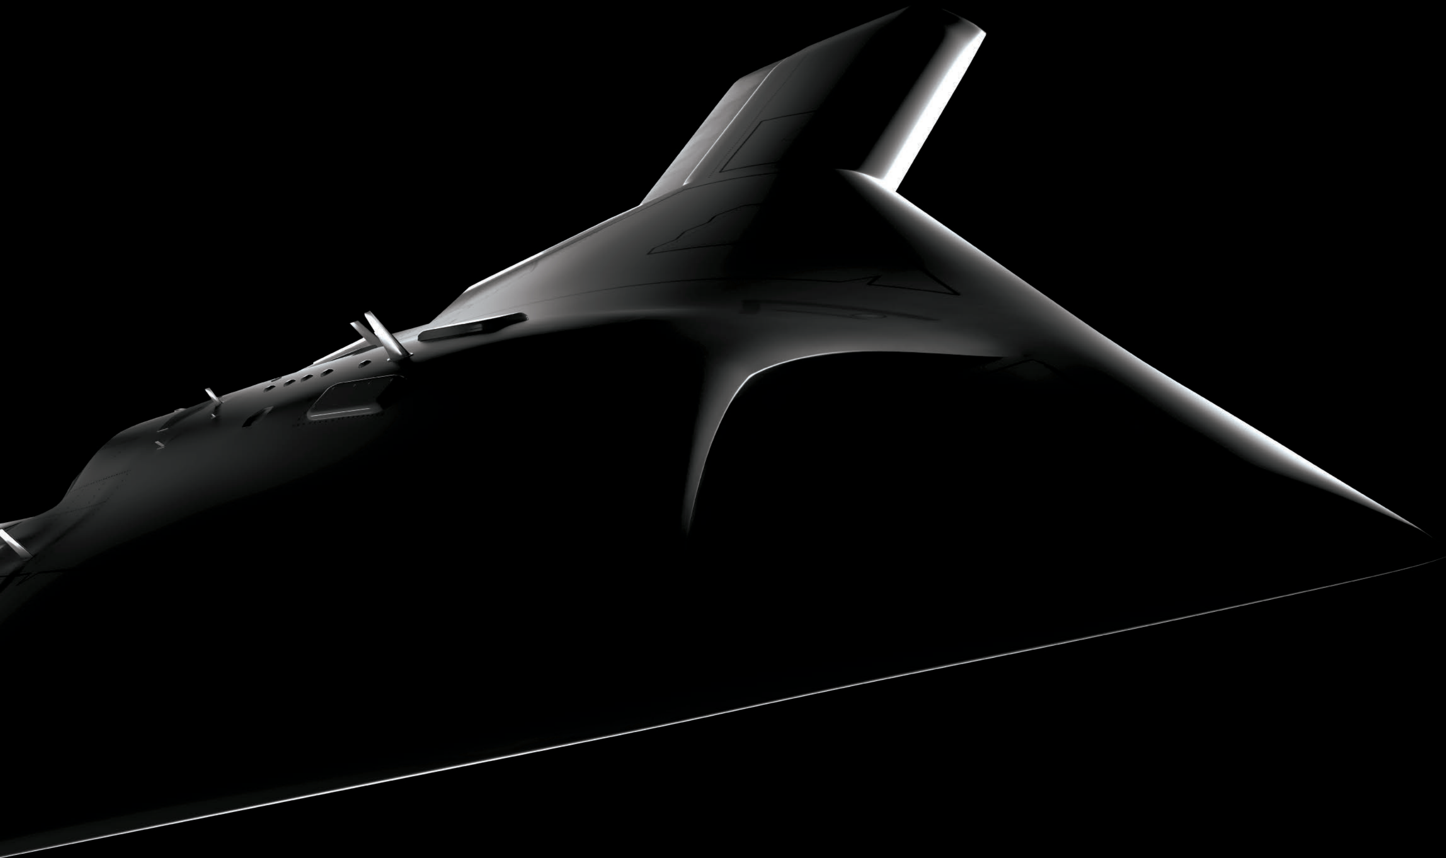

From the Editor

Three years ago, when CRISPR-Cas9, a method of precisely editing DNA, suddenly appeared liked a technology from the future, scientists realized we could now engineer the human race by giving our children and their heirs improved genes. But many disavowed that we would do so soon.

The biologist Weizhi Ji, who created two gene-edited macaque monkeys at Kunming Biomedical International, said that creating humans with CRISPR-edited genomes was "very possible," but added that "considering the safety issue, there would still be a long way to go." (See "10 Breakthrough Technologies: Genome Editing," May/June 2014.)

Yet I was sure that scientists would research how to edit the human germ line, and quickly. ("Germ line" is biologists' jargon for the egg and the sperm, which combine to form an embryo. Editing the DNA of such cells, or of the embryo itself, would pass heritable changes to future generations.)

They would try because editing genes with CRISPR was trivially easy, and it would be a sensational thing to do. "Any scientist with molecular biology skills and knowledge of how to work with [embryos] is going to be able to do this," says Jennifer Doudna, a biologist at the University of California, Berkeley, who in 2012 co-discovered how to use CRISPR to edit genes.

Mostly, they would research CRISPR because it seems a powerful way to prevent disease from birth. Guoping Feng, a neurobiologist at MIT's McGovern Institute for Brain Research, believes that gene-edited human beings are "10 to 20 years away," but nonetheless approves of human germ-line editing. Feng says, "To me, it's possible in the long run to dramatically improve health, lower costs. It's a kind of prevention."

Why not use CRISPR to eliminate diseases like Huntington's, a terrible, fatal neurodegenerative disorder triggered by a defect in a single gene? Or why not correct the DNA of an embryo with a mutation in a gene called *BRCA1*, which causes ovarian and breast cancer? While you're fiddling with an embryo's DNA, why not insert naturally occurring gene variants that confer extraordinary characteristics like unbreakable bones or resistance to diseases like Alzheimer's?

As our biomedicine editor, Antonio Regalado, reports in this issue's cover story, "Engineering the Perfect Baby" (page 26), experiments designed to correct the DNA in a woman's egg or a man's sperm, or to directly edit the DNA of an early-stage embryo using CRISPR, are already being carried out.

Why not? One concern is that the technologies would not be widely available, at least at first. Their expense would mean only rich people would have perfect children. Another worry is that germ-line engineering would affect unborn people without their consent. The most potent objection is that we don't know what we're doing: if you provide immunity to a disease, you might break something in a genome.

History insists that when a technology has obvious utility, it will be used. But *how* we use a new technology is our choice. In March, writing in the journal *Science*, a group of scientists including Doudna and two Nobel laureates called for a great debate on the genetic engineering of humans and a moratorium on any effort to create engineered babies. Crucially, the scientists did not ask their peers to stop using CRISPR to edit human embryos for research purposes; but they recommended convening a "globally representative" group of government agencies, ethics experts, and scientists to suggest policies to guide that research.

Let's have that debate. Write to me at jason.pontin@technologyreview.com.

GUIDO VITTI

Contents

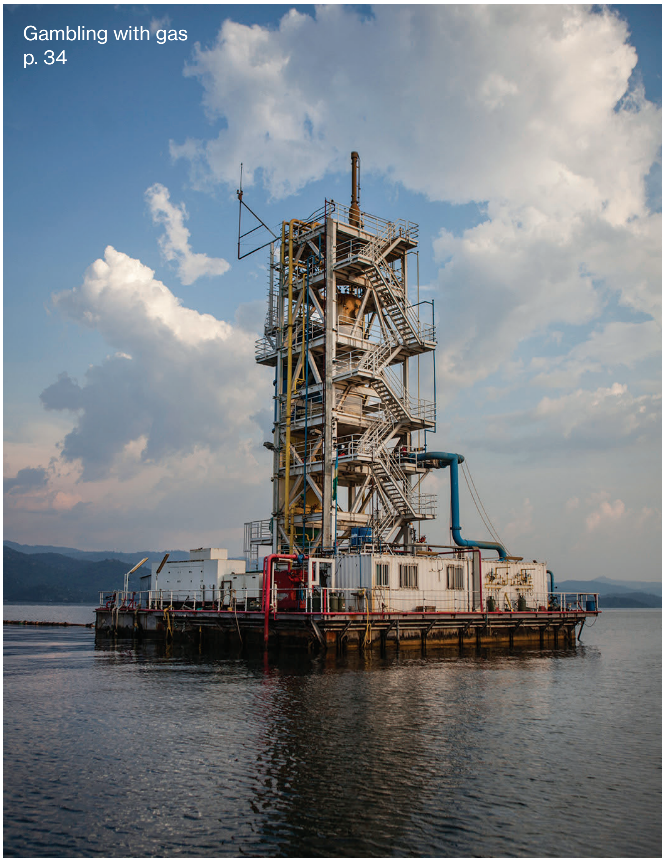

Gambling with gas
p. 34

PHOTOGRAPH BY JASON FLORIO

May/June 2015

ON THE COVER:

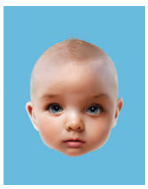

Illustration by Arn0

MIT TECHNOLOGY REVIEW
VOL. 118 | NO. 3

TECHNOLOGYREVIEW.COM

CUSTOMER SERVICE AND SUBSCRIPTION INQUIRIES

National: 800-877-5230

International: 903-636-1115

E-mail: technologyreview@pubservice.com

Web: www.technologyreview.com/customerservice

MIT Records: 617-253-8270 (alums only)

Reprints:
techreview@wrightsmedia.com
877-652-5295

Licensing and permissions:
licensing@technologyreview.com

Technology Review
One Main Street, 13th Floor
Cambridge, MA 02142
Tel: 617-475-8000

The mission of *MIT Technology Review* is to equip its audiences with the intelligence to understand a world shaped by technology.

Technology Review, Inc., is an independent nonprofit 501(c)(3) corporation wholly owned by MIT; the views expressed in our publications and at our events are not always shared by the Institute.

De Technologia non multum scimus. Scimus autem, quid nobis placeat.

Feedback

E-mail letters@technologyreview.com

Write *MIT Technology Review*
One Main Street, 13th Floor
Cambridge, MA 02142

Please include your address, telephone number, and e-mail address.

Letters and comments may be edited for both clarity and length.

Five Most Popular Stories

MIT Technology Review
Volume 118, Number 2

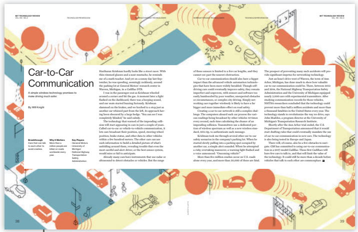

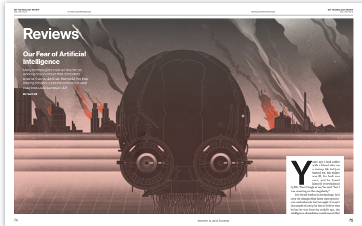

1 10 Breakthrough Technologies: Magic Leap

The real value of Magic Leap's platform is its potential light field display, combined with very complicated algorithms. We don't know whether Magic Leap has solved these problems or not—but we do know their demo and technical feasibility plan were solid enough to convince investors to put down half a billion dollars.
—Juan David Hincapié-Ramos

This is just another reason to detach from reality and focus on someone else's version of reality.
—psenatori

2 10 Breakthrough Technologies: Car-to-Car Communication

If traffic lights communicate to let cars adjust their speed to arrive only at green lights, everyone is going to want this.
—Gary Kreie

The auto industry has a long way to go to understand the hazards of connected cars, as shown by researchers hacking Tesla and BMW. Those hacks were just a nuisance, like unlocking the doors, but what happens when people start trusting warnings from other cars approaching, when that information has been manipulated?
—apptimates

3 Why We Don't Have Battery Breakthroughs

It is not realistic to expect "breakthroughs" in any area of technology. They rarely occur, and are almost never the product of efforts directed at a specific result. Instead, battery technology will advance at a slow, steady pace, improving a little year after year.
—big.league.slider

I'm an electric-car advocate but I can see that giving up the ability to drive as far as you want is enough to make people think twice. Fast recharge will fix the problem faster than increased range—for me, at least. **—asdar**

4 10 Breakthrough Technologies: Internet of DNA

We need better models of consent for the sharing of genomic data, but there won't be one solution that works globally: each country will have to develop its own response, guided by the work at the international level.
—Chris Arnold, chairman Human Variome Project International

Most privacy concerns are not rational assessments of the real risk of your information being processed. I would want everyone's DNA databased if it meant I could spend five more years with my parents. **—orangesherbet**

5 Our Fear of Artificial Intelligence

The purpose of technology is not to enable human ability. The purpose is to make money. We shouldn't fear intelligent bots, but we should be deathly afraid of those who make and sell them. **—ar18**

The easiest pathway to super-intelligence, whole-brain emulation, does not require any theoretical breakthroughs, just better technology than we now have. I see no value, given the potential ramifications of such technology, in pretending that there is nothing to worry about.
—HockeyAndReason

The Perils of Genome Sharing

Genetic privacy is a hugely complicated problem ("10 Breakthrough Technologies: Internet of DNA," March/April 2015), but *MIT Technology Review* treats it as some sort of side note or inconvenience in the way of the latest gee-whiz technical development.

Disclosing your genome opens up you and your entire extended family to life-altering discrimination. Today, these risks are somewhat obscure, but in a few years, they will be all too obvious as research uncovers more and more risks associated with various genetic markers. So people who disclose their genomes based on current risk-reward calculations have in fact made a very bad mistake, not just on their own behalf but also on behalf of their yet-unborn children. —mrnorwood

People who disclose their genomes based on current risk-reward calculations have made a very bad mistake, not just on their own behalf but on behalf of their yet-unborn children.

The Computer in Your Car

Computers in vehicles can have unintended consequences ("10 Breakthrough Technologies: Car-to-Car Communication"). My son drives long haul and for a couple of years drove a new truck with a braking system that would stop the truck if it detected an imminent threat. Sounds like a good idea. But the braking system would stop the truck on a clear road with no traffic or any obstacle around (maybe it saw a ghost?). Going around a sharp turn with fencing along the curve, the radar sensed an object in front, made no allowances for the wheels being turned, and stopped the truck dead in its tracks. There were numerous other incidents—most at highway speeds.

—NJdriver

It's Too Late to Stop AI

Regarding "Our Fear of Artificial Intelligence" (March/April 2015): the genie is out of the bottle. It's too late to put it back in, and why would we want to anyway? Life has gotten so much better for people around the globe in the past 200 years thanks in part to the brilliance and persistence of researchers, inventors, and businesspeople doing often lonely and heroic work against impossible odds to drive us all forward. Will their work lead to an unintended consequence like the extinction of humanity? Well, the preservation of our species is largely out of our control—it's always been like that and always will be. —Keith Kulper

Bitcoin Boondoggle

Bitcoin has now been around for a considerable period of time ("A Weekend in Bitcoin City," March/April 2015). Doesn't the fact that articles like this are being written signify that Bitcoin hasn't succeeded as a payment medium? Can people live for a weekend on cash? Can people live for a weekend on a credit or debit card? Of course they can. Bitcoin offers nothing new in this regard. And since it offers no payment dispute or fraud protections as commercial payment methods do, it offers a lot less. Bitcoin is a technological bust. Why can't people accept that? —TAG

Correction

In the March/April issue, we incorrectly credited a photo of Steven Chu to Imke Lass/Redux. The photo should have been credited to Charles Ommanney/Getty Images.

Views

Jennifer A. Doudna

Daniel Kammen

Martha Kim

BIOTECHNOLOGY

Genome Gambits

New genetic tools can do tremendous good—if we use them carefully.

My dad loved to hike in the rain forests near our home on the Big Island of Hawaii, often to hunt for mushrooms with Don Hemmes, his colleague at the University of Hawaii. The goal of these trips was not to harvest mushrooms but to photograph them for a research project that Hemmes was leading. When I accompanied them, I was always struck by the incredible diversity of the mushrooms we found. Having learned a little about genetics in school, I wondered what kinds of DNA changes were responsible for these organisms' range of colors, shapes, and sizes. And how could we figure out such molecular signatures?

Fast-forward 30-odd years, and it's become routine to sequence the entire genomes of organisms, and to interpret that information to reveal the underlying causes of observable traits. A simple and effective technology for making precise changes to those genomic sequences, developed by harnessing a system that bacteria use to fight viral infections, has exploded into widespread use (see "Engineering the Perfect Baby," page 26). The technology, called CRISPR, relies on a programmable DNA-cutting enzyme called Cas9, together with its guide RNA, to let scientists alter the genetic information within cells, tissues, and whole organisms. Scientists have used it to generate new strains of wheat, to cure a genetic disease in the livers of adult mice, and to produce altered fungal cells capable of efficient biofuel synthesis. The CRISPR-Cas9 technology has opened up a world of research opportunities that were inconceivable just three years ago. The technology will benefit humanity in many ways.

There's also a growing appreciation of the risks involved. CRISPR-Cas9 technology can, as an example, be used to alter the DNA in germ cells or embryos, resulting in permanent changes to the genetic makeup of every differentiated cell in a resulting organism—and to that organism's progeny. The system is so efficient that genetic changes it introduces could become self-propagating. Such applications could be employed to cure genetic disease in humans or to limit the fitness of disease-carrying organisms—but the intricacies of genetic interaction mean those uses could also have unintended consequences, perhaps triggering other diseases.

Research is needed to understand the utility and risks of CRISPR-Cas9 in cells including human germ cells, as well as the risks inherent in any human clinical applications that might be possible in the future. We should research the ramifications of using genome engineering to control organisms, such as mosquitoes, that can spread malaria or dengue fever. While we should embrace this technology, scientists also must come together to guide peers and regulators as to its responsible use.

Jennifer A. Doudna is a professor of biology and chemistry at the University of California, Berkeley. She was one of the inventors of the CRISPR technology.

ENERGY

Peace Through Grids

How smart energy policy can ease conflicts.

When I served as the chief technical specialist for renewable energy and energy efficiency at the World Bank, one project I found especially interesting was the construction of an electricity highway between the rich geothermal energy

ANDY FRIEDMAN

fields of the Rift Valley in Kenya through the Lake Turkana plains—where the best wind resource identified to date in Africa was recently mapped—to newly constructed hydroelectric facilities in Ethiopia. Not only are these indigenous renewable energy resources largely untapped, but the policy tools to build markets for clean energy are often most effective in poor, power-starved nations.

Roughly 1.5 billion people around the world live without electricity today, so these kinds of projects should be a priority for international development (see "Lake Kivu's Great Gas Gamble," page 34). But such projects have ramifications well beyond energy. They represent a major opportunity to use some of our greatest infrastructure investments to build peaceful, prosperous, and coöperative regional economies. The grid is the greatest engineering achievement of the 20th century, and yet we've given very little thought to building partnerships through shared energy commerce. This has to change.

Critical opportunities now exist to build coöperative regional economies and address the global climate crisis. One example can be found in South Sudan, the world's newest nation, where old ethnic tensions, exacerbated by potential oil and gas wealth, have disrupted an already fragile process of nation building. But if investors could connect South Sudan to the emerging Eastern Africa Power Pool, the country could disengage from its tense relationship with Sudan and cheaply power the local economy—in a place where less than 2 percent of the population now has electricity.

Kosovo, one of Europe's poorest nations, has been a battleground over a proposed coal-fired power plant. Kosovo happens to have significant resources in wind, biomass, and hydropower, much of which would most efficiently be developed jointly with Albania. This approach would make the coal plant—a pollution-belcher six kilometers from the capital city—an unnecessary anachronism. Kosovo and Albania recently announced that they will integrate their power markets, a step that could unleash the region's impressive solar energy resources to work closely with bioenergy and distributed hydropower.

Nations linked by energy commerce—and in particular clean, local energy—are far less likely to enter into hostilities than those that see each other only as regional rivals. That's why governments seeking to build strong international partnerships would do well to make transmission diplomacy and development a centerpiece of foreign policy. Such efforts would greatly aid energy access globally and make clean energy the technology of choice for a new wave of investments.

Daniel Kammen is a professor of energy at the University of California, Berkeley, and a contributing lead author for the Intergovernmental Panel on Climate Change.

COMPUTING

Better Architecture

Computers are overdue for the fundamental changes they could soon get.

Computer architectures aren't laws of physics. They're man-made inventions designed to harness raw resources, such as billions of transistors, for a range of useful computational tasks.

When our computing needs and tasks change—as they inevitably will over the decades—it becomes increasingly awkward to express programs through the original architecture. And yet that's where we find ourselves—adhering to an ossified architecture that imposes constraints and slows our technological progress.

Today's architectures are more than half a century old. In the 1940s, electronic computers became reprogrammable, with data and instructions (a.k.a. software) stored in memory and passed to a central processing unit (CPU) for computation. This architecture evolved slightly over time but remained fundamentally the same.

The vast majority of computing devices today are connected to the Internet, making them vulnerable to remote attack. Our data centers demand the type of strong security—including isolation and tracking of data—that classic architectures were never designed to support.

That's one reason computing architectures must evolve. A system being developed by Hewlett-Packard, known as the Machine (see "Machine Dreams," page 48), uses electronic components called memristors to store and process information—offering more powerful ways to handle large amounts of data—together with silicon photonic components that allow data to be transported at very high speeds using light. HP's researchers are also developing a new operating system, Machine OS, to make the most of this new architecture.

Reinvention like this doesn't solve all our problems. In some cases it creates new ones. The consistent architecture of IBM's System 360 in the 1960s and 1970s ensured that buyers of early models could upgrade their machines and feel confident that the programs they were already using would continue to work. Can a new architecture evolve without forcing every program to evolve with it?

Probably. Since the days of the System 360, compilers and program translators—tools that allow software to run on different architectures—have matured substantially. We'll need to make the most of such tools if we hope to loosen our ties to legacy architectures and allow computers like the Machine to thrive.

Martha Kim is an associate professor of computer science at Columbia University.

Upfront

An Oak Ridge engineer inspects a molten-salt reactor in a shipping vessel in 1964.

COURTESY OF OAK RIDGE NATIONAL LABORATORY

A Blast from the Nuclear Past

A new version of an old reactor design could help make nuclear power safer and more economical.

A new take on an old reactor design could make nuclear power cleaner and safer, and therefore more competitive with fossil fuels.

Terrestrial Energy, a startup based in Ontario, Canada, is commercializing the reactor design, which is based on work done at Oak Ridge National Laboratory

Upfront

in Tennessee. Terrestrial plans to start licensing the design in Canada later this year. It is designing a reactor that uses molten salt rather than water as a coolant, which essentially makes the design meltdown-proof.

Researchers at Oak Ridge first demonstrated a molten-salt reactor in the 1960s, and they have tested various designs over the past several decades. Terrestrial has modified one of these designs in ways it says will make the technology cheap enough to deploy.

Nuclear power could play an important role in replacing fossil fuels. However, conventional nuclear reactors cost far more to build than fossil-fuel power plants, in large part because of safety regulations stipulating costly redundant pumps, containment structures, and other components designed to prevent a catastrophic accident. Terrestrial CEO Simon Irish says the company's new molten-salt design could make it possible to simplify and reduce the cost of safety systems.

In a molten-salt nuclear reactor, if the power goes off or the reactor is damaged, the system will cool off on its own without allowing radioactivity to spread. Conventional nuclear reactors must be actively cooled with water continuously pumped through them. If the pumps stop, the fuel will start to overheat, and this can lead to a meltdown and release radioactive materials into the environment.

A handful of other startups, including Transatomic Power, are working to commercialize molten-salt reactors as well. The technology is also a focus of considerable R&D efforts in China and elsewhere around the world.

Terrestrial's designs are more conventional than those being developed by Transatomic. The company plans to use the same materials employed in reactors tested at Oak Ridge, whereas Transatomic's designs use several new materials.

Terrestrial is designing a reactor that uses molten salt rather than water as a coolant.

In Terrestrial's reactor, uranium is mixed with a molten-salt coolant. If the fuel gets too hot, that causes the mixture to expand, which slows down fission and reduces the heat of the fuel. This automatically regulates the temperature and prevents overheating. Also, the coolant boils only at a very high temperature, so unlike water, it won't evaporate even if the pumps stop working. Furthermore, if the reactor was damaged and the fuel and coolant mixture leaked out, the fission reactions would slow down and the molten fuel would solidify, limiting the spread of radioactive material.

Irish says the design will also reduce nuclear waste by about two-thirds because the reactor operates at temperatures twice as hot as those in a conventional reactor, which improves its efficiency and reduces the amount of fuel needed. In addition, he says that recycling the fuel, which further reduces waste, is simpler than it is with conventional reactors.

To make the Oak Ridge design more practical, Terrestrial modified it so that the reactor could be manufactured in a factory and shipped to a plant site on the back of a truck. Another distinctive feature of the Terrestrial design is that key parts are disposable. One challenge with the Oak Ridge molten-salt design is that a critical material, graphite, doesn't last very long, which would mean a plant's operators would have to replace it fairly regularly. The new design houses the reactor's main components, including the graphite, in a sealed unit that can be swapped out every seven years, theoretically making operating the plant easier.

Terrestrial has produced a preliminary design and is working with Oak Ridge on a more detailed design, which an engineering firm could use to produce blueprints. The company hopes to see the first commercial reactor started in the early part of the next decade. *—Kevin Bullis*

TO MARKET

Holographics

Leia

COMPANY:
Leia Display

PRICE:
To be announced

AVAILABILITY:
Late 2015

A company called Leia has developed a small display module capable of producing full-color 3-D images and videos that are visible—with no special glasses—from 64 different viewpoints. Inventor David Fattal came up with the idea while working on optical interconnects at HP Labs. Optical interconnects rely on nanoscale structures called diffraction gratings, which cause the light rays that hit them to travel in precise directions. Fattal realized that the idea could also be used to display holographic images, and he left Hewlett-Packard to pursue that idea with his own company. Leia has developed a way to generate the holograms using a conventional liquid crystal display (LCD).

COURTESY OF LEIA

Banking for Beginners

A nonprofit experiments with using digital currency to give poor teenagers their first experience with financial services.

Proponents of the digital currency Bitcoin have often argued that money made out of computer code could help poor people access financial services. But so far applications for the technology have been almost exclusively aimed at people with Internet access and smartphones. Now a South African nonprofit is preparing to give the idea its first real test.

In the next few months, some teenage girls in poor areas of South Africa will be offered the chance to test out a kind of digital savings account operated via text message. It will be offered as a new feature on a mobile social-network service the girls already use. The savings feature will let people earn and save mobile airtime credit, which is used in addition to government-backed money as a currency in some countries.

Behind the scenes, the new service is powered by a digital currency called Stellar, which was inspired by Bitcoin. Stellar, like Bitcoin, is based on a system that uses cryptographic software to create digital tokens that can't be counterfeited. But Stellar differs from Bitcoin in that it is designed to act as an intermediary between conventional currencies and assets, to speed up transfers between them, and not as a means of payment in its own right. Development of Stellar is being undertaken by a nonprofit, the Stellar Development Foundation, backed by $3 million from the payments company Stripe.

The Praekelt Foundation develops a piece of free software called Vumi, which powers interactive services that can run over text messages on phones without data plans. Humanitarian organizations including UNICEF, USAID, and the Gates Foundation use Vumi to deliver health and education programs in Africa and elsewhere. The new savings feature will be offered as an opt-in addition to exist-

PING ZHU

ing social-networking services built on Vumi and aimed at teenage girls living in poverty. The feature works by rewarding teenagers with small amounts of airtime in return for sending messages, reading things, and participating in other activities on the social service.

Tests of the savings service are planned for several countries. Kentaro Toyama, an associate professor at the University of Michigan who studies technology and development (see Q&A, page 24), says that Stellar will still need to win the approval of regulators. Most countries' financial rules make it difficult for companies that are not banks to transfer and store money, often for good reason, says Toyama. That can be a problem for organizations trying something new, like offering savings or transfer services on mobile devices. *—Tom Simonite*

QUOTED

"The dream of package deliveries by drone is ahead of its time."

—Phil Finnegan, an analyst with the Teal Group, an aerospace consultancy. He argues that technological obstacles hampered Amazon's and Google's plans for automated package delivery even before U.S. regulators proposed strict regulations on drones.

"A watch is a very covert object. I could see a new kind of private language or low-level communication emerging from this kind of wearable, using pulses or squeezes."

—Laura Seargeant Richardson, a user-experience expert at Argodesign, speaking about the Apple Watch.

"I am actually incredibly afraid of what happens if Internet.org succeeds. We need to be behind the Internet as an open marketplace for ideas."

—Mark Surman, executive director of the Mozilla Foundation, discussing Facebook's efforts to subsidize mobile Internet services, and access to its services, in the developing world.

BY THE NUMBERS

1,000,000

The number of people whose genomes will be sequenced under the U.S. government's "precision medicine" initiative.

80

The percentage of Internet content that is only available in one of 10 languages, according to the World Bank.

Upfront

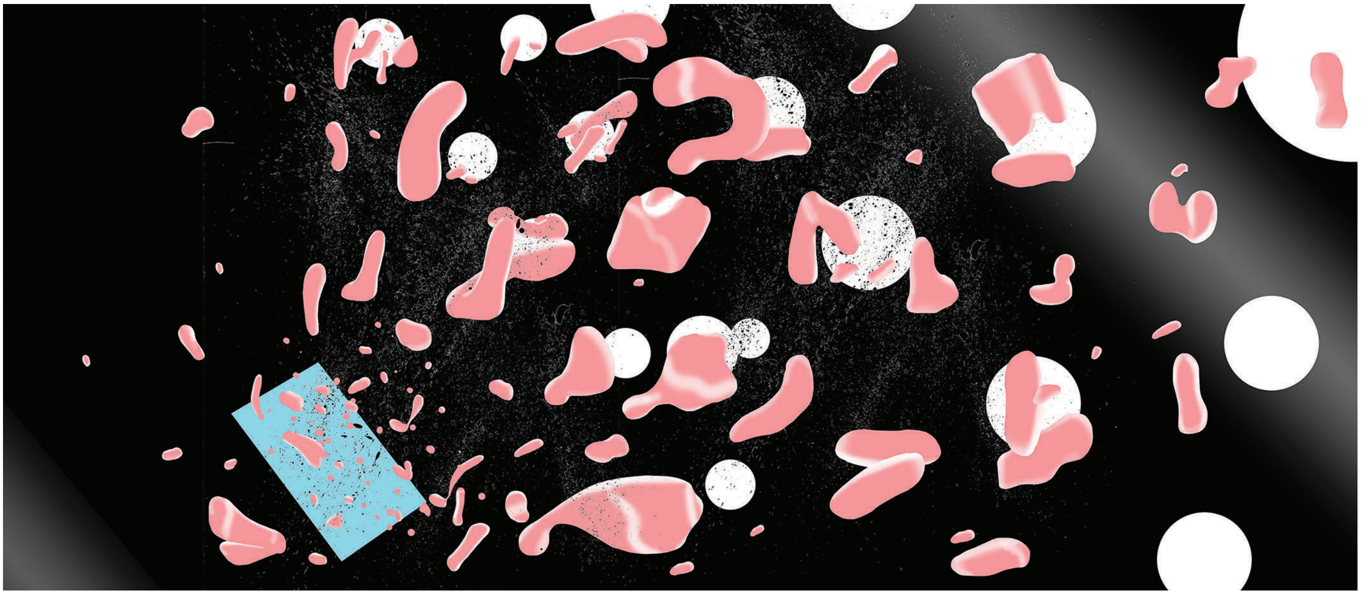

JON HAN

Beating Diabetes

Pancreatic tissue grown from stem cells is being implanted in some patients.

Fourteen years ago, during the darkest moments of the "stem-cell wars" pitting American scientists against the White House of George W. Bush, one group could be counted on to urge research using cells from human embryos: parents of children with type 1 diabetes. Motivated by scientists who told them these cells would lead to amazing cures, they spent millions on TV ads, lobbying, and countless phone calls to Congress.

Now the first test of a type 1 diabetes treatment using stem cells has finally begun. In October, a San Diego man had two pouches of lab-grown pancreas cells, derived from human embryonic stem cells, inserted into his body through incisions in his back. Two other patients have since received the stand-in pancreas, from a San Diego company called ViaCyte.

It's a significant step, partly because the ViaCyte study is only the third in the United States of any treatment based on embryonic stem cells. These cells, once removed from early-stage human embryos, can be grown in a lab dish and retain the ability to differentiate into any cell or tissue type.

Type 1 patients must constantly monitor their blood glucose using finger pricks, carefully time when and what they eat, and routinely inject themselves with insulin. This hormone triggers the removal of excess glucose from the blood. In type 1 diabetics, the pancreas doesn't make it because their own immune system has destroyed the pancreatic islets, the clusters of cells containing the insulin-secreting beta cells.

Despite years of research, there is still "just nothing" to offer patients, says Robert Henry, a doctor at the University of California, San Diego, whose center is carrying out the surgeries for ViaCyte.

Henry slightly overstates the case, but not by much. There is something called the Edmonton Protocol, a surgical technique that uses islets collected from cadavers. Early hopes for the Edmonton Protocol were quickly tempered, however. It requires recipients to take powerful immune-suppressing drugs for life, and suitable donor pancreases are in extremely short supply.

The protocol's initial success came only two years after the discovery of embryonic stem cells, in 1998. Those pressing for a diabetes cure quickly set a new goal: pair something like the Edmonton Protocol with the technology of lab-grown beta cells.

"It was obvious that if we had another cell source that was replenishable, large numbers [of people] would benefit," says Richard Insel, chief scientific officer of the Juvenile Diabetes Research Foundation (JDRF), a nonprofit with 300,000 members.

That's why the JDRF battled the restrictions threatened by the Bush White House, and why its members were behind a 2004 voter initiative in California that created the California Institute


HOW DO YOU KNOW WHAT YOUR CUSTOMERS WANT WHEN THEY DON’T EVEN WANT IT YET?
IT’S SIMPLE.
THE ANSWER IS SAP HANA.
What if you could predict the future by mastering the present? Major retailers use SAP HANA® to analyze hundreds of millions of transactions from tens of millions of customers to predict trends years in advance.
For more, go to sap.com/simple
© 2015 SAP SE. All rights reserved. Results based on specific system configurations. Customer results may vary.
SAP
Run Simple

Upfront

for Regenerative Medicine, a state agency authorized to spend $3 billion on stem-cell research. The California institute has given ViaCyte six grants worth $39 million, more than it's given any other company, and JDRF has invested another $14 million directly.

Although the idea of growing replacement beta cells is conceptually simple, in practice it's proved more difficult to execute than anyone imagined. One challenge has been getting stem cells to turn into real, functioning pancreas cells, especially the insulin-secreting beta cells. ViaCyte's approach is to grow immature pancreas cells, counting on the body to do the work of transforming them into actual beta cells. The second problem is how to evade a patient's immune system, which will attack any transplanted cell. ViaCyte's solution is a plastic mesh capsule, which it fills with about 40 million of the immature pancreas cells. The purpose of the capsule is to screen out the immune system's killer T cells while allowing the transplanted cells to receive nourishment from the bloodstream, as well as to sense blood sugar and respond.

Data that ViaCyte supplied to the U.S. Food and Drug Administration last year in order to receive approval for the human trial showed that the cells produced insulin, glucagon (secreted in response to low blood sugar), and somatostatin, a growth hormone, and successfully regulated blood sugar, at least in mice.

Though the current human trial is meant mostly to test for safety, Henry suspects that his patients may see some reduction in their need for injected insulin. From the first patient, whose identity hasn't been disclosed, Henry says he has already retrieved a test sack, which he says appeared to be functioning properly. Nobody is sure how long the implanted cells will survive, but it is certain that patients would have to have new implants installed periodically.

ViaCyte has shown it can regulate diabetes by implanting a capsule filled with cells, at least in mice.

At least two other groups say they've also controlled diabetes in rodents and may soon start trials of their own. One is BetaLogics Venture, a subsidiary of the drug giant Johnson & Johnson, which last year reported reversing diabetes in mice using what its patents describe as a yarn-based scaffold in a polyester shell.

Douglas Melton, a biologist at Harvard University who has two children with type 1 diabetes, worries that the ViaCyte system may not work. He thinks deposits of fibrotic, scarlike tissue will glom onto the capsules, starving the cells inside of oxygen and blocking their ability to sense sugar and release insulin. Melton also thinks it might take immature cells up to three months to become fully functional. And many won't become beta cells, winding up as other types of pancreatic cells instead.

Last October, Melton's group announced it had managed to grow fully mature, functional beta cells in the lab, a scientific first. Melton thinks implanting mature cells would allow a bioartificial pancreas to start working right away.

To encapsulate his cells, Melton has been working with bioengineer Daniel Anderson at MIT to develop their own capsule. Anderson doesn't want to say exactly how it works, but a recent patent filing from his lab describes a container made of layers of hydrogels, some containing cells and others anti-inflammatory drugs.

After the stem-cell wars, and then a decade of trying to turn the technology's promises into reality, Henry says he feels convinced that "cells in bags" of some kind are going to be the answer to type 1 diabetes. He's aware that curing rodents doesn't guarantee a good result for people, but he says the clinical trial is another in a series of "small steps" toward much-improved lives for millions. "I am just so positive that this is the future," he says.

—Brian Alexander

TO MARKET

Speedy Printing

CLIP

COMPANY:
Carbon3D

PRICE:
To be announced

AVAILABILITY:
To be announced

LARS SAHL

A new additive-manufacturing technology is 25 to 100 times faster than conventional 3-D printing and also produces stronger parts at a lower cost. The technology was developed by a startup called Carbon3D, which was spun out of the University of North Carolina at Chapel Hill. Conventional 3-D printers produce objects by depositing or solidifying one layer of material at a time, a process that lends itself to making very complex objects but is much slower and more expensive than large-volume manufacturing techniques. Carbon3D's technology makes objects continuously rather than in discrete layers, and it is fast enough to potentially compete with conventional mass manufacturing.

Upfront

Cheaper, Better Batteries Ahead

A promising new battery technology is nearing commercialization thanks to backing from a major appliance manufacturer.

A new kind of lithium-ion battery that stores twice as much energy as anything on the market today is a step closer to commercialization, thanks to a deal between Sakti3, the startup behind the technology, and Dyson, a British company that makes vacuum cleaners, fans, and other devices for the home.

Sakti3 and Dyson have signed a joint development agreement, and Dyson plans to incorporate the startup's batteries into some of its future products. It hasn't said what those products might be, but cordless power tools and vacuum cleaners would appear to be likely candidates. Dyson also announced a $15 million investment in Sakti3.

Sakti3 is developing what is known as a solid-state battery, meaning it replaces the liquids used in a conventional lithium-ion battery with solid materials. This makes it possible to use electrode materials that store more energy. For these reasons, solid-state batteries are widely regarded as one of the most promising types of advanced battery technology out there. But such batteries have proved difficult to manufacture, and they usually need to be very thin because current does not flow easily through a solid electrolyte material.

Sakti3 has developed a type of solid-state battery than can be made with inexpensive manufacturing technology and can also store even more energy. Though it's not disclosing many details about how it's accomplished this, it has relied on highly sophisticated and proprietary computer models to design its batteries. It makes the batteries using novel materials and manufacturing techniques.

Sakti3's design also does away with the flammable liquid electrolyte used in conventional lithium-ion batteries, and this makes it feasible to use a different set of high-energy storage materials.

If the battery technology can be commercialized successfully, it could go on to have applications not only in consumer electronics but also in electric vehicles, the range of which is currently limited by the capacity of conventional batteries.

Of course, many battery startups have struggled to translate impressive lab results into products, in part because their prototypes are custom-made, often using expensive manufacturing techniques that can't easily be scaled up. Sakti3 makes its prototypes on standard manufacturing equipment, and this could help ensure that the technology, once perfected, can be commercialized.

The first solid-state lithium-ion batteries may be used in cordless power tools and vacuum cleaners.

Success isn't a sure thing, though. Making a few prototypes might not translate to high yields in a factory, where conditions and raw materials can vary. But the partnership with Dyson, which has extensive manufacturing experience and relatively deep pockets, could make the transition easier for Sakti3.

In a statement, Dyson founder and chief engineer James Dyson said, "Sakti3 has achieved leaps in performance which current battery technology simply can't."

—Kevin Bullis

3 QUESTIONS

Andrew Mason

After being fired from Groupon, the first company you founded, you've launched Detour, which offers audio tours. How can it stand out from other tour apps?

If you've ever used other audio-tour apps, or if you've gone to a museum, they all have this similar problem: you're constantly fidgeting with your phone. For the walking-tour apps, you might be looking at a map and clicking on pins on the map and then playing content or something like that. We wanted something where people could just have an experience that feels like you're there with a member of the community or whoever it is, and the technology just melts out of the way.

I was impressed with how well the audio matched up with my walking speed—can you explain the technology behind it?

In the script you'll put markers or cues for when you want the listener to start walking, and then you know that you have a certain amount of time until they get to the next stop, and you want the pacing to make sense, and then there's testing around noise levels in the environment that can be confounders.

What are you doing in places where GPS doesn't work very well?

We're launching an architecture tour later this month that's in a high-rise zone, financial district, in San Francisco. And GPS gets way off in those zones, so we need something that's a little bit more reliable. We're setting up iBeacons along the way that will create a better read. And then we're using other sensors on the phone.

—Rachel Metz

PING ZHU

Augmented Advertising

If virtual and augmented reality take off, expect the technologies to become a new platform for commercials.

I'm sitting in a desk chair in an office in Mountain View, California. But with a virtual-reality headset strapped to my head and headphones over my ears, it looks and sounds as if I'm standing in the belly of a blimp, flying high above silent city blocks dotted with billboards for a *Despicable Me* theme-park ride.

The blimp ride is part of a demo built by MediaSpike, a startup that's making ads for virtual reality. Even the blimp itself is an ad: before boarding it, I can see that its exterior is covered with a larger-than-life version of one of the film's short, yellow characters.

For now, augmented reality and virtual reality are not widely used. But as new headsets hit the market, advertisers will surely stake out virtual ground. Companies like Facebook-owned Oculus VR, Sony, Microsoft, and Magic Leap are working on consumer hardware, some of which is slated to begin selling late this year. Already, Oculus VR and Samsung have released the Gear VR, a $199 developer-geared headset that lets you experience 3-D games and videos by inserting a Samsung smartphone into the device.

MediaSpike is among several companies aiming to serve brands that want to advertise on these new platforms. Founder and CEO Blake Commagere says the company started out a few years ago working to bring sponsored content to smartphone- and tablet-based games. Now the company is thinking about how billboards, videos, and other kinds of product placement can fit into the computer-generated worlds viewed on devices like the Gear VR, as well as on headsets that don't yet have a firm release date.

BENDIK KALTENBORN

From inside the headgear, it still looks pretty primitive. Before the blimp ride, I was driving around MediaSpike's digitally rendered town, which is empty of activity beyond the movie billboards, the blimp, and a giant display in a vacant town square that's showing the *Minions* trailer. I found it nauseating, too (a common complaint from people using stereoscopic 3-D technologies, because of the disconnect between visual and physical senses). But it's still a lot more interesting than the banner and pop-up ads we're used to seeing on websites and mobile apps.

The huge value of the online advertising market suggests it could be lucrative to experiment in this area. According to Magna Global, a media market researcher and investor, digital-media revenue rose 17 percent in 2014, to $142 billion. It's expected to climb another 15 percent, to $163 billion globally, this year.

One company that has already had some success merging virtual ads with the real world is Blippar. You can take a smartphone app created by the company and use it to view sponsored augmented-reality content, such as virtual football players on Pepsi cans, or virtual nail polish shades from Maybelline.

There will be plenty of challenges ahead for companies working to bring ads to virtual and augmented reality. Beyond the obvious difficulty of working with emerging technologies, Hart and others are going to need to figure out what kinds of ads people will respond positively to, and which they might find too overwhelming or obnoxious.

Commagere acknowledges that just as with product placement in movies, it will be possible to overdo it with virtual products. (One example of this phenomenon occurred during my blimp ride, when I turned around to see a pixelated can of Pepsi perched next to the craft's control panel.) But he's hopeful that this will feel just a little less intrusive than conventional pop-ups and other Web ads.

—Rachel Metz

Upfront

Power to the People

The relationship between electricity and economic development (see "Lake Kivu's Great Gas Gamble," page 34) is clear in India, where the poorest states have the least access to the grid. Nonetheless, while electricity is usually necessary for prosperity, it's not sufficient.

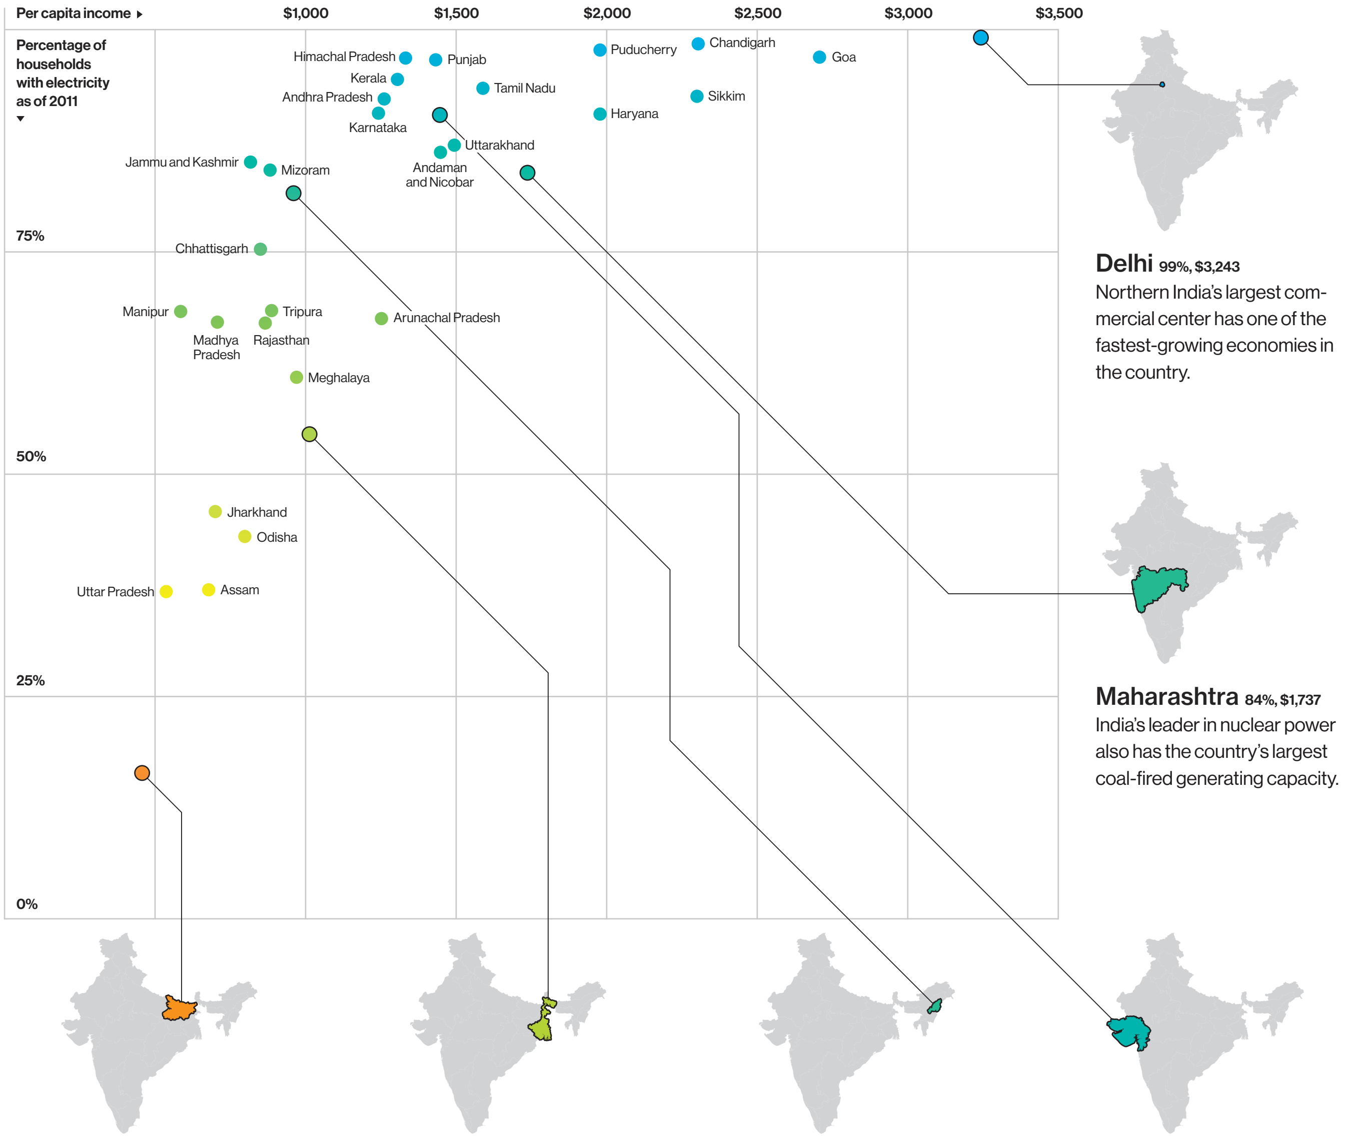

Delhi 99%, $3,243
Northern India's largest commercial center has one of the fastest-growing economies in the country.

Maharashtra 84%, $1,737
India's leader in nuclear power also has the country's largest coal-fired generating capacity.

Bihar 16%, $457
The third-most-populous Indian state has added very little power capacity in the past decade.

West Bengal 55%, $1,013
Despite lacking widespread access to the grid, India's fourth-most-populous state is a big agricultural producer.

Nagaland 82%, $960
Though access is high in this agricultural state, electricity supply often falls short.

Gujarat 90%, $1,446
This state is India's leading producer of solar power and will soon be home to the nation's biggest solar-panel factory.

ILLUSTRATION BY LUKE SHUMAN
DATA FROM GOVERNMENT OF INDIA;
INCOME DATA WAS UNAVAILABLE FOR THREE UNION TERRITORIES THAT ARE NOT LISTED

SARAH MAZZETTI

Rewriting the "Imitation Game"

Some computer scientists are searching for more meaningful ways to measure artificial intelligence.

We have self-driving cars, knowledgeable digital assistants, and software capable of putting names to faces as well as any expert. But do these displays of machine aptitude represent genuine intelligence? For decades artificial-intelligence experts have struggled to find a practical way to answer the question.

"Asking whether an artificial entity is 'intelligent' is fraught with difficulties," says Mark Riedl, an associate professor at Georgia Tech. "Eventually a self-driving car will outperform human drivers, so we can even say that along one dimension, an AI is super-intelligent. But we might also say that it is an idiot savant, because it cannot do anything else, like recite a poem or solve an algebra problem."

The most famous effort to measure machine intelligence does not resolve the challenges involved; instead, it obscures them. In his 1950 paper "Computing Machinery and Intelligence," the British computer scientist Alan Turing considered the capacity of computers and turned to a black-box definition: if we accept humans as an intelligent species, then anything that exhibits behaviors indistinguishable from human behavior must also be intelligent.

Turing also proposed a test, called the "imitation game," in which a computer would prove its intelligence by convincing a person, through conversation, that it was also human. In the years since, the Turing test has been widely adopted and also widely criticized—not because of flaws in Turing's original idea, but because of flaws in its execution.

A chatbot called Eugene Goostman made headlines last June for passing the Turing test in a contest organized at the University of Reading in the U.K. The software convinced 30 percent of the human judges involved that it was human. But the chatbot relies on obfuscation and subterfuge rather than the natural back and forth of intelligent conversation.

"Turing's original description mandated a freewheeling conversation that could range over any subject, and there was no nonsense allowed," says Leora Morgenstern, an expert on AI who works at Leidos, a defense contractor headquartered in Virginia. AI's earliest proponents hoped to work toward some form of general intelligence. But as the complexity of the task unfurled, research fractured into smaller, more manageable tasks. This produced progress, but it also turned machine intelligence into something that could not easily be compared with human intellect.

Most AI researchers still pursue highly specialized areas, but some are turning their attention back to generalized intelligence and considering new ways to measure progress. For Morgenstern, a machine will demonstrate intelligence only when it can show that once it knows one intellectually challenging task, it can easily learn another related task. Riedl agrees that the test should be broad: "Conversation is just one aspect of human intelligence. Creativity is another. Problem solving and knowledge are others."

Riedl has designed an alternative to the Turing test, which he has dubbed the Lovelace 2.0 test (a reference to Ada Lovelace, a 19th-century English mathematician who programmed a seminal calculating machine). Riedl's test would focus on creative intelligence, with a human judge challenging a computer to create something—a story, poem, or drawing—and gradually making the task harder.

This test might not be the ideal successor to the Turing test. But it may offer a better way to understand machine intelligence than any simple test. "Who is to say being above a certain score is intelligent or being below is unintelligent?" Riedl says. "Would we ever ask such a question of humans?" *—Simon Parkin*

Q+A

PATRICK LEGER

Reality Check

Kentaro Toyama calls himself "a recovering technoholic"—someone who once was "addicted to a technological way of solving problems." Five years in India changed him. After getting his PhD in computer science and working on machine vision technologies at Microsoft, Toyama moved to Bangalore in 2004 to help lead the company's new research center there. He and his colleagues launched dozens of projects that sought to use computers and Internet connectivity to improve education and reduce poverty. But early successes in pilot projects often couldn't be replicated; in some schools, computers made things worse. In a book being released this spring, *Geek Heresy: Rescuing Social Change from the Cult of Technology*, Toyama argues that technologists undermine efforts at social progress by promoting "packaged interventions" at the expense of more difficult reforms. Toyama, who is now an associate professor in the School of Information at the University of Michigan, spoke to *MIT Technology Review*'s deputy editor, Brian Bergstein.

When you went to India, technological optimism was flourishing there. Bangalore was being described as the next great tech hub.
Yes, absolutely. We [reasoned] that here's this technology sector, which is incredibly successful. Isn't there some way that we can take the sheen of this sector and spread it around not just to people who are well educated and middle class, but also to people who are poorer and who don't have the same kind of educational advantages—basically, to the 80 percent of the country that really is considered poor by any standard? At that time, there were barely even mobile phones. It was mostly Internet-connected PCs. I thought there was some way to use them in a way that we could support the health-care system, agriculture, or education.

What was your success rate?
I ultimately took stock of 50-odd projects that I had either been directly involved with or supervised. Very few were the kind where we felt, "This is working so well that we should really expand it." Very often, it was because there were just limits to the human and institutional capacity on the ground that could take advantage of the technology.

For example, in education, one of the most difficult things to overcome is the way in which education is done—everything from how the public school system is managed to how it's administered to how the government interacts with it. In India, we found instances where teachers were often called away by the government. The government feels that they're government employees and, therefore, can be called upon to help with other government tasks.

Another example is the health-care system. If you go to a typical rural clinic, it's not the kind of place that anybody from the United States would think of as a decent place to get health care. Bringing along a laptop, connecting it to wireless,

and providing Internet so you can do telemedicine is just an incredibly thin cover. It's a thin, superficial change.

As an example of a project that has made a difference, your book cites Digital Green. It makes and shows videos in which farmers in India share advice about planting techniques or how to handle animals. What makes that successful?

We were very cautious that the technology doesn't replace an existing agriculture extension system. It merely amplifies whatever system is already there: human beings who have a lot of agriculture expertise, who are willing to talk to farmers and who have some rapport with farmers.

We do these sessions in villages where somebody who is in touch with an agriculture extension person will call together the villagers and then do a screening using small projectors. On the one hand, it's just basically a video screening. [But] the mediators are trained in a way so that they're asked to provoke discussion, which is a critical part of the learning process. If you don't do the mediation, it's just like watching TV. And the farmers, many of them have TVs in their homes. They see agricultural programs, but that information does not register. They don't end up implementing it for various reasons. Whereas when they have discussions together, or when they see farmers that are just like them, then they're much more likely to believe the content of the videos and adopt the farming practice.

What do you think of One Laptop per Child, which may be the poster child for the idea of a technologically driven intervention?

There are already several randomized, controlled trials of schools with and without One Laptop per Child. Generally, what most of these studies show is that schools with laptops did not see their children gain anything in terms of academic achievement, in terms of grades, in terms of test scores, in terms of attendance, or in terms of supposed engagement with the classroom.

That might surprise people who have seen anecdotal success stories.

It's the anecdotes that really keep the technology sector going in the [economic] development context. It is so easy to get an interesting story if you take some gadget and give it to a child. I have done this myself multiple times. The first thing that you see is kids just overjoyed that they have this new gadget in their hands. It's a new toy, and they love it. You can't not take a photograph of a smiling kid holding a laptop.

The reality is, that joy is the same joy that you see when you peek over the shoulder of a kid who has a smartphone in their hands in the developed world, which is to say they're overjoyed because they're playing Angry Birds. On the one hand, I do think that a certain amount of educational toys and play is important, but I just don't think that K through 12 education of any serious kind can be based entirely on that kind of play.

> **"The fundamental error people make is that we should have the computer be the primary instrument of education. It's not clear to me why people seem to make that leap."**

I think it's perfectly sensible for parents to want a certain amount of exposure to technology for their children, both as a form of explorative play and as a way to get them used to technology that they'll undoubtedly encounter later in their life. I think the fundamental error people make is that, therefore, we should have the computer be the primary instrument of education for all children. That, to me, is a major leap. It's not clear to me why people seem to make that leap all the time. I think one of the issues is we tend to think of education as being the content. We overemphasize the importance of content, as opposed to emphasizing the part that's really difficult in any good education, which is adult-supervised motivation—the motivation of the child to learn something.

Why do many technologists fail to see that technology on its own is so limited in the changes it can bring?

What I see is a societal level of confusion of correlation and cause. We see this incredible success of Silicon Valley and the technology industry overall. On a daily basis, especially those of us who can afford the technology, we see it in our own lives: here's this technology that just seems to be making everything more convenient, everything better, and so on and so forth. So we assume that it's the technology that is directly responsible, when in fact it's a whole bunch of other stuff that already has to be there in the first place. If you're lacking that other stuff to begin with, then the technology by itself doesn't cause all of those benefits.

The tricky thing about this is you can be very scientific about these things and still come up with the wrong conclusion. Multiple times in my lab, we'd run trials where you compare a control situation with a treatment situation. The treatment situation gets some kind of technology. If you measure some positive benefit in the technology case, your conclusion is that technology helped. But it was always the people that we worked with, the partners that we chose and the people on the ground who interacted with the people that we wanted to support. All of those human factors were required for the technology itself to have an impact; whether the technology helped or not was really up to people.

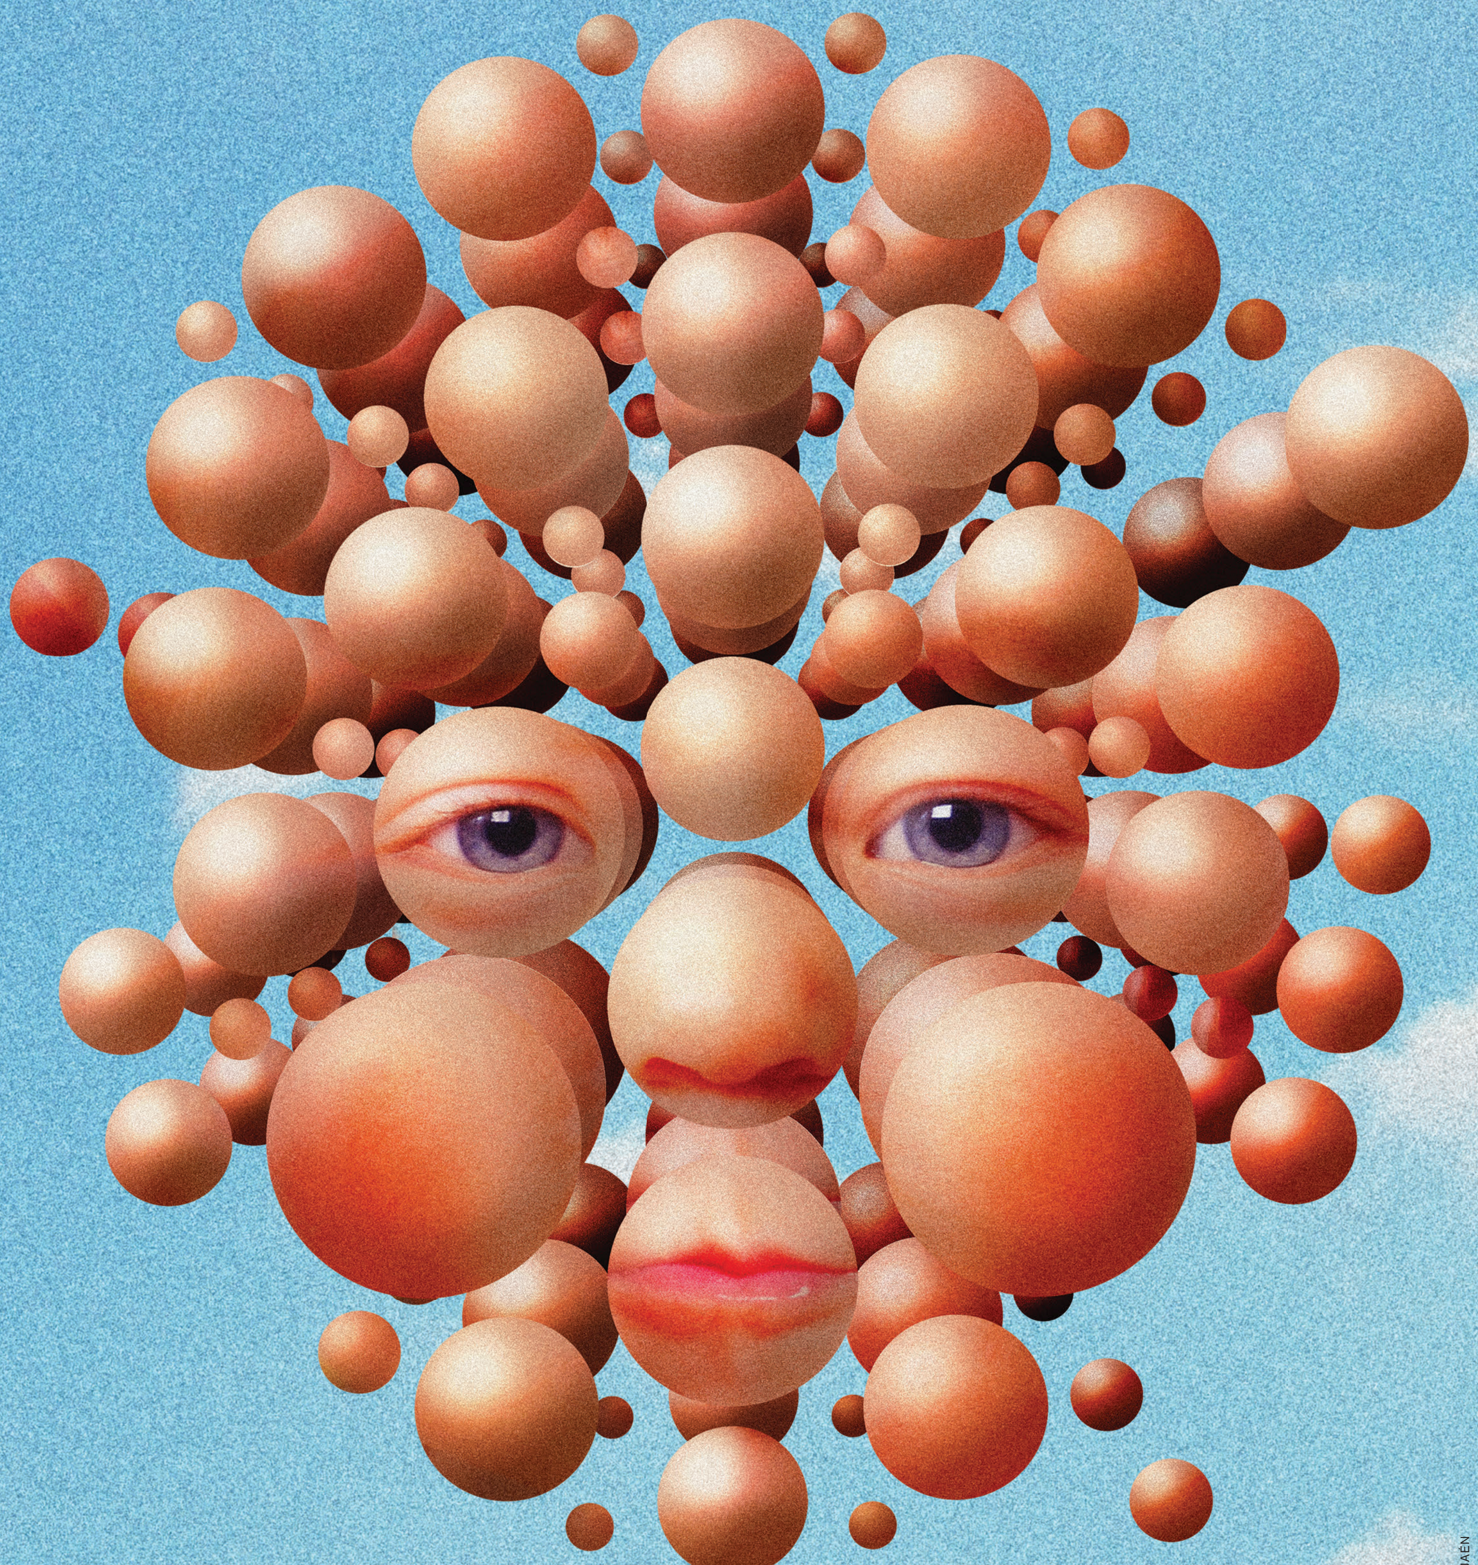

JAVIER JAÉN

Engineering the Perfect Baby

Scientists are developing ways to edit the DNA of tomorrow's children. Should they stop before it's too late?

By Antonio Regalado

If anyone had devised a way to create a genetically engineered baby, I figured George Church would know about it.

At his labyrinthine laboratory on the Harvard Medical School campus, you can find researchers giving *E. Coli* a novel genetic code never seen in nature. Around another bend, others are carrying out a plan to use DNA engineering to resurrect the woolly mammoth. His lab, Church likes to say, is the center of a new technological genesis—one in which man rebuilds creation to suit himself.

When I visited the lab last June, Church proposed that I speak to a young postdoctoral scientist named Luhan Yang. A Harvard recruit from Beijing, she'd been a key player in developing a powerful new technology for editing DNA, called CRISPR-Cas9. With Church, Yang had founded a small biotechnology company to engineer the genomes of pigs and cattle, sliding in beneficial genes and editing away bad ones.

As I listened to Yang, I waited for a chance to ask my real questions: Can any of this be done to human beings? Can we improve the human gene pool? The position of much of mainstream science has been that such meddling would be unsafe, irresponsible, and even impossible. But Yang didn't hesitate. Yes, of course, she said. In fact, the Harvard laboratory had a project under way to determine how it could be achieved. She flipped open her laptop to a PowerPoint slide titled "Germline Editing Meeting."

Here it was: a technical proposal to alter human heredity. "Germ line" is biologists' jargon for the egg and sperm, which combine to form an embryo. By editing the DNA of these cells or the embryo itself, it could be possible to correct disease genes and pass those genetic fixes on to future generations. Such a technology could be used to rid families of scourges like cystic fibrosis. It might also be possible to install genes that offer lifelong protection against infection, Alzheimer's, and, Yang told me, maybe the effects of aging. Such history-making medical advances could be as important to this century as vaccines were to the last.

That's the promise. The fear is that germ-line engineering is a path toward a dystopia of superpeople and designer babies for those who can afford it. Want a child with blue eyes and blond hair? Why not design a highly intelligent group of people who could be tomorrow's leaders and scientists?

Just three years after its initial development, CRISPR technology is already widely used by biologists as a kind of search-and-replace tool to alter DNA, even down to the level of a single letter. It's so precise that it's expected to turn into a promising new approach for gene therapy in people with devastating illnesses. The idea is that physicians could directly correct a faulty gene, say, in the blood cells of a patient with sickle-cell anemia (see "Genome Surgery," March/April 2014). But that kind of gene ther-

apy wouldn't affect germ cells, and the changes in the DNA wouldn't get passed to future generations.

In contrast, the genetic changes created by germ-line engineering would be passed on, and that's what has made the idea seem so objectionable. So far, caution and ethical concerns have had the upper hand. A dozen countries, not including the United States, have banned germ-line engineering, and scientific societies have unanimously concluded that it would be too risky to do. The European Union's convention on human rights and biomedicine says tampering with the gene pool would be a crime against "human dignity" and human rights.

But all these declarations were made before it was actually feasible to precisely engineer the germ line. Now, with CRISPR, it is possible.

The experiment Yang described, though not simple, would go like this: The researchers hoped to obtain, from a hospital in New York, the ovaries of a woman undergoing surgery for ovarian cancer caused by a mutation in a gene called *BRCA1*. Working with another Harvard laboratory, that of antiaging specialist David Sinclair, they would extract immature egg cells that could be coaxed to grow and divide in the laboratory. Yang would use CRISPR in these cells to correct the DNA of the *BRCA1* gene. They would try to create a viable egg without the genetic error that caused the woman's cancer.

Yang would later tell me that she dropped out of the project not long after we spoke. Yet it remained difficult to know if the experiment she described was occurring, canceled, or awaiting publication. Sinclair said that a collaboration between the two labs was ongoing, but then, like several other scientists whom I'd asked about germ-line engineering, he stopped replying to my e-mails.

Regardless of the fate of that particular experiment, human germ-line engineering has become a burgeoning research concept. At least three other centers in the United States are working on it, as are scientists in China, in the U.K., and at a biotechnology company called OvaScience, based in Cambridge, Massachusetts, that boasts some of the world's leading fertility doctors on its advisory board.

The objective of these groups is to demonstrate that it's possible to produce children free of specific genes involved in inherited disease. If it's possible to correct the DNA in a woman's egg, or a man's sperm, those cells could be used in an in vitro fertilization (IVF) clinic to produce an embryo and then a child. It might also be possible to directly edit the DNA of an early-stage IVF embryo using CRISPR. Several people interviewed by *MIT Technology Review* said that such experiments had already been carried out in China and that results describing edited embryos were pending publication. These people, including two high-ranking specialists, didn't wish to comment publicly because the papers are under review.

All this means that germ-line engineering is much further along than anyone imagined. "What you are talking about is a major issue for all humanity," says Merle Berger, one of the founders of Boston IVF, a network of fertility clinics that is among the largest in the world and helps more than a thousand women get pregnant each year. "It would be the biggest thing that ever happened in our field." Berger predicts that repairing genes involved in serious inherited diseases will win wide public acceptance but says the idea of using the technology beyond that would cause a public uproar because "everyone would want the perfect child": people might pick and choose eye color and eventually intelligence. "These are things we talk about all the time," he says. "But we have never had the opportunity to do it."

Editing embryos

How easy would it be to edit a human embryo using CRISPR? Very easy, experts say. "Any scientist with molecular biology skills and knowledge of how to work with [embryos] is going to be able to do this," says Jennifer Doudna, a biologist at the University of California, Berkeley, who in 2012 co-discovered how to use CRISPR to edit genes.

To find out how it could be done, I visited the lab of Guoping Feng, a biologist at MIT's McGovern Institute for Brain Research, where a colony of marmoset monkeys is being established with the aim of using CRISPR to create accurate models of human brain diseases. To create the models, Feng will edit the DNA of embryos and then transfer them into female marmosets to produce live monkeys. One gene Feng hopes to alter in the animals is *SHANK3*. The gene is involved in how neurons communicate; when it's damaged in children, it is known to cause autism.

Feng said that before CRISPR, it was not possible to introduce precise changes into a primate's DNA. With CRISPR, the technique should be relatively straightforward. The CRISPR

All this means that germ-line engineering is much further along than anyone imagined.

Genetic Modification of Babies

Percentage of U.S. adults saying that changing a baby's genetic characteristics for each purpose is ...

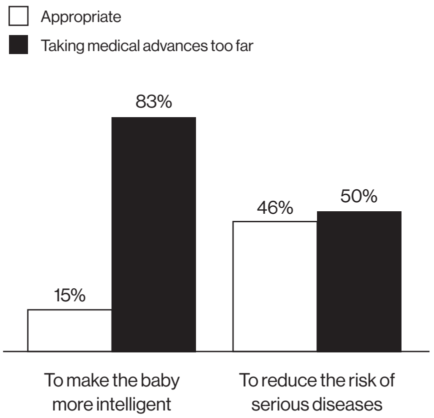

SURVEY OF U.S. ADULTS AUGUST 15-25, 2014, BY PEW RESEARCH CENTER. THOSE SAYING "DON'T KNOW" ARE NOT SHOWN.

system includes a gene-snipping enzyme and a guide molecule that can be programmed to target unique combinations of the DNA letters, A, G, C, and T; get these ingredients into a cell and they will cut and modify the genome at the targeted sites.

But CRISPR is not perfect—and it would be a very haphazard way to edit human embryos, as Feng's efforts to create gene-edited marmosets show. To employ the CRISPR system in the monkeys, his students simply inject the chemicals into a fertilized egg, which is known as a zygote—the stage just before it starts dividing.

Feng said the efficiency with which CRISPR can delete or disable a gene in a zygote is about 40 percent, whereas making specific edits, or swapping DNA letters, works less frequently—more like 20 percent of the time. Like a person, a monkey has two copies of most genes, one from each parent. Sometimes both copies get edited, but sometimes just one does, or neither. Only about half the embryos will lead to live births, and of those that do, many could contain a mixture of cells with edited DNA and without. If you add up the odds, you find you'd need to edit 20 embryos to get a live monkey with the version you want.

That's not an insurmountable problem for Feng, since the MIT breeding colony will give him access to many monkey eggs and he'll be able to generate many embryos. However, it would present obvious problems in humans. Putting the ingredients of CRISPR into a human embryo would be scientifically trivial. But it wouldn't be practical for much just yet. This is one reason that many scientists view such an experiment (whether or not it has really occurred in China) with scorn, seeing it more as a provocative bid to grab attention than as real science. Rudolf Jaenisch, an MIT biologist who works across the street from Feng and who in the 1970s created the first gene-modified mice, calls attempts to edit human embryos "totally premature." He says he hopes these papers will be rejected and not published. "It's just a sensational thing that will stir things up," says Jaenisch. "We know it's possible, but is it of practical use? I kind of doubt it."

For his part, Feng told me he approves of the idea of germ-line engineering. Isn't the goal of medicine to reduce suffering? Considering the state of the technology, however, he thinks actual gene-edited humans are "10 to 20 years away." Among other problems, CRISPR can introduce off-target effects or change bits of the genome far from where scientists had intended. Any human embryo altered with CRISPR today would carry the risk that its genome had been changed in unexpected ways. But, Feng said, such problems may eventually be ironed out, and edited people will be born. "To me, it's possible in the long run to dramatically improve health, lower costs. It's a kind of prevention," he said. "It's hard to predict the future, but correcting disease risks is definitely a possibility and should be supported. I think it will be a reality."

Editing eggs

Elsewhere in the Boston area, scientists are exploring a different approach to engineering the germ line, one that is technically more demanding but probably more powerful. This strategy combines CRISPR with unfolding discoveries related to stem cells. Scientists at several centers, including Church's, think they will soon be able to use stem cells to produce eggs and sperm in the laboratory. Unlike embryos, stem cells can be grown and multiplied. Thus they could offer a vastly improved way to create edited offspring with CRISPR. The recipe goes like this: First, edit the genes of the stem cells. Second, turn them into an egg or sperm. Third, produce an offspring.

asdfasdf

CRISPR in the Human Gene Pool
Key players in the development of human germ-line editing.

Jennifer Doudna
University of California, Berkeley
Key achievement: In 2012, co-developed CRISPR editing system, using bacteria.
Current work: Has expressed concern about CRISPR technology and its potential germ-line applications.

George Church
Harvard Medical School
Key achievement: In 2013, demonstrated that CRISPR can work in human cells.
Current work: Project to engineer genomes of animals, including pigs whose organs can be transplanted to human patients.

OvaScience
Cambridge, Massachusetts
Key achievement: Raised nearly $300 million to commercialize egg stem-cell technology for IVF clinics.
Current work: Exploring the possibility of editing eggs to remove disease genes from future generations.

Jinsong Li
Shanghai Institute for Biological Sciences
Key achievement: In 2015, corrected a genetic disease in mice by editing the DNA of sperm cells.
Current work: Hopes to edit human sperm and demonstrate medical applications for IVF procedures.

Xingxu Huang
ShanghaiTech University
Key achievement: In 2014, was part of team that bred the first CRISPR-edited monkeys in China.
Current work: Seeking permission to genetically modify discarded IVF embryos.

Azim Surani
University of Cambridge, U.K.
Key achievement: In late 2014, showed that human skin cells can be turned into primitive eggs and sperm.
Current work: Using CRISPR in stem cells to study basic questions in developmental biology.

Some investors got an early view of the technique on December 17, at the Benjamin Hotel in Manhattan, during commercial presentations by OvaScience. The company, which was founded four years ago, aims to commercialize the scientific work of David Sinclair, who is based at Harvard, and Jonathan Tilly, an expert on egg stem cells and the chairman of the biology department at Northeastern University (see “10 Emerging Technologies: Egg Stem Cells,” May/June 2012). It made the presentations as part of a successful effort to raise $132 million in new capital during January.

During the meeting, Sinclair, a velvet-voiced Australian whom *Time* last year named one of the “100 Most Influential People in the World,” took the podium and provided Wall Street with a peek at what he called “truly world-changing” developments. People would look back at this moment in time and recognize it as a new chapter in “how humans control their bodies,” he said, because it would let parents determine “when and how they have children and how healthy those children are actually going to be.”

The company has not perfected its stem-cell technology—it has not reported that the eggs it grows in the lab are viable—but Sinclair predicted that functional eggs were “a when, and not an if.” Once the technology works, he said, infertile women will be able to produce hundreds of eggs, and maybe hundreds of embryos. Using DNA sequencing to analyze their genes, they could pick among them for the healthiest ones.

Genetically improved children may also be possible. Sinclair told the investors that he was trying to alter the DNA of these egg stem cells using gene editing, work he later told me he was doing with Church’s lab. “We think the new technologies with genome editing will allow it to be used on individuals who aren’t just interested in using IVF to have children but have healthier children as well, if there is a genetic disease in their family,” Sinclair told the investors. He gave the example of Huntington’s disease, caused by a gene that will trigger a fatal brain condition even in someone who inherits only one copy. Sinclair said gene editing could be used to remove the lethal gene defect from an egg cell. His goal, and that of OvaScience, is to “correct those mutations before we generate your child,” he said. “It’s still experimental, but there is no reason to expect it won’t be possible in coming years.”

Sinclair spoke to me briefly on the phone while he was navigating in a cab across a snowed-in Boston, but later he referred my questions to OvaScience. When I contacted OvaScience, Cara Mayfield, a spokeswoman, said its executives could not comment because of their travel schedules but confirmed that the company was working on treating inherited disorders with gene editing. What was surprising to me

was that OvaScience's research in "crossing the germ line," as critics of human engineering sometimes put it, has generated scarcely any notice. In December of 2013, OvaScience even announced it was putting $1.5 million into a joint venture with a synthetic-biology company called Intrexon, whose R&D objectives include gene-editing eggs to "prevent the propagation" of human disease "in future generations."

When I reached Tilly at Northeastern, he laughed when I told him what I was calling about. "It's going to be a hot-button issue," he said. Tilly also said his lab was trying to edit egg stem cells with CRISPR "right now" to rid them of an inherited genetic disease that he didn't want to name. Tilly emphasized that there are "two pieces of the puzzle"—one being stem cells and the other gene editing. The ability to create large numbers of egg stem cells is critical, because only with sizable quantities can genetic changes be stably introduced using CRISPR, characterized using DNA sequencing, and carefully studied to check for mistakes before producing an egg.

Tilly predicted that the whole end-to-end technology—cells to stem cells, stem cells to sperm or egg and then to offspring—would end up being worked out first in animals, such as cattle, either by his lab or by companies such as eGenesis, the spinoff from the Church lab working on livestock. But he isn't sure what the next step should be with edited human eggs. You wouldn't want to fertilize one "willy nilly," he said. You'd be making a potential human being. And doing that would raise questions he's not sure he can answer. He told me, "'Can you do it?' is one thing. If you can, then the most important questions come up. 'Would you do it? Why would you want to do it? What is the purpose?' As scientists we want to know if it's feasible, but then we get into the bigger questions, and it's not a science question—it's a society question."

Improving humans

If germ-line engineering becomes part of medical practice, it could lead to transformative changes in human well-being, with consequences to people's life span, identity, and economic output. But it would create ethical dilemmas and social challenges. What if these improvements were available only to the richest societies, or the richest people? An in vitro fertility procedure costs about $20,000 in the United States. Add genetic testing and egg donation or a surrogate mother, and the price soars toward $100,000.

Others believe the idea is dubious because it's not medically necessary. Hank Greely, a lawyer and ethicist at Stanford University, says proponents "can't really say what it is good for." The problem, says Greely, is that it's already possible to test the DNA of IVF embryos and pick healthy ones, a process that adds about $4,000 to the cost of a fertility procedure. A man with Huntington's, for instance, could have his sperm used to fertilize a dozen of his partner's eggs. Half those embryos would not have the Huntington's gene, and those could be used to begin a pregnancy.

Indeed, some people are adamant that germ-line engineering is being pushed ahead with "false arguments." That is the view of Edward Lanphier, CEO of Sangamo Biosciences, a California biotechnology company that is using another gene-editing technique, called zinc fingers nucleases, to try to treat HIV in adults by altering their blood cells. "We've looked at [germ-line engineering] for a disease rationale, and there is none," he says. "You can do it. But there really isn't a medical reason. People say, well, we don't want children born with this, or born with that—but it's a completely false argument and a slippery slope toward much more unacceptable uses."

The fear? A dystopia of superpeople and designer babies for those who can afford it.

Critics cite a host of fears. Children would be the subject of experiments. Parents would be influenced by genetic advertising from IVF clinics. Germ-line engineering would encourage the spread of allegedly superior traits. And it would affect people not yet born, without their being able to agree to it. The American Medical Association, for instance, holds that germ-line engineering shouldn't be done "at this time" because it "affects the welfare of future generations" and could cause "unpredictable and irreversible results." But like a lot of official statements that forbid changing the genome, the AMA's, which was last updated in 1996, predates today's technology. "A lot of people just agreed to these statements," says Greely. "It wasn't hard to renounce something that you couldn't do."

Others predict that hard-to-oppose medical uses will be identified. A couple with several genetic diseases at once might not be able to find a suitable embryo. Treating infertility is another possibility. Some men don't produce any sperm, a condition called azoospermia. One cause is a genetic defect in which a region of about one million to six million DNA let-

ters is missing from the Y chromosome. It might be possible to take a skin cell from such a man, turn it into a stem cell, repair the DNA, and then make sperm, says Werner Neuhausser, a young Austrian doctor who splits his time between the Boston IVF fertility-clinic network and Harvard's Stem Cell Institute. "That will change medicine forever, right? You could cure infertility, that is for sure," he says.

I spoke with Church several times by telephone over the last few months, and he told me what's driving everything is the "incredible specificity" of CRISPR. Although not all the details have been worked out, he thinks the technology could replace DNA letters essentially without side effects. He says this is what makes it "tempting to use." Church says his laboratory is focused mostly on experiments in engineering animals. He added that his lab would not make or edit human embryos, calling such a step "not our style."

What is Church's style is human enhancement. And he's been making a broad case that CRISPR can do more than eliminate disease genes. It can lead to augmentation. At meetings, some involving groups of "transhumanists" interested in next steps for human evolution, Church likes to show a slide on which he lists naturally occurring variants of around 10 genes that, when people are born with them, confer extraordinary qualities or resistance to disease. One makes your bones so hard they'll break a surgical drill. Another drastically cuts the risk of heart attacks. And a variant of the gene for the amyloid precursor protein, or APP, was found by Icelandic researchers to protect against Alzheimer's. People with it never get dementia and remain sharp into old age.

Church thinks CRISPR could be used to provide people with favorable versions of genes, making DNA edits that would act as vaccines against some of the most common diseases we face today. Although he told me anything "edgy" should be done only to adults who can consent, it's obvious to him that the earlier such interventions occur, the better.

Church tends to dodge questions about genetically modified babies. The idea of improving the human species has always had "enormously bad press," he wrote in the introduction to *Regenesis*, his 2012 book on synthetic biology, whose cover was a painting by Eustache Le Sueur of a bearded God creating the world. But that's ultimately what he's suggesting: enhancements in the form of protective genes. "An argument will be made that the ultimate prevention is that the earlier you go, the better the prevention," he told an audience at MIT's Media Lab last spring. "I do think it's the ultimate preventive, *if* we get to the point where it's very inexpensive, extremely safe, and very predictable." Church, who has a less cautious side, proceeded to tell the audience that he thought changing genes "is going to get to the point where it's like you are doing the equivalent of cosmetic surgery."

Some thinkers have concluded that we should not pass up the chance to make improvements to our species. "The human genome is not perfect," says John Harris, a bioethicist at Manchester University, in the U.K. "It's ethically imperative to positively support this technology." By some measures, U.S. public opinion is not particularly negative toward the idea. A Pew Research survey carried out last August found that 46 percent of adults approved of genetic modification of babies to reduce the risk of serious diseases.

The same survey found that 83 percent said genetic modification to make a baby smarter would be "taking medical advances too far." But other observers say higher IQ is exactly what we should be considering. Nick Bostrom, an Oxford philosopher best known for his 2014 book *Superintelligence*, which raised alarms about the risks of artificial intelligence in computers, has also looked at whether humans could use reproductive technology to improve human intellect. Although the ways in which genes affect intelligence aren't well understood and there are far too many relevant genes to permit easy engineering, such realities don't dim speculation on the possibility of high-tech eugenics.

What if everyone could be a little bit smarter? Or a few people could be a lot smarter? Even a small number of "super-enhanced" individuals, Bostrom wrote in a 2013 paper, could change the world through their creativity and discoveries, and through innovations that everyone else would use. In his view, genetic enhancement is an important long-range issue like climate change or financial planning by nations, "since human problem-solving ability is a factor in every challenge we face."

To some scientists, the explosive advance of genetics and biotech means germ-line engineering is inevitable. Of course, safety questions would be paramount. Before there's a genetically edited baby saying "Mama," there would have to be

> "The human genome is not perfect. It's ethically imperative to positively support this technology."

tests in rats, rabbits, and probably monkeys, to make sure they are normal. But ultimately, if the benefits seem to outweigh the risks, medicine would take the chance. “It was the same with IVF when it first happened,” says Neuhausser. “We never really knew if that baby was going to be healthy at 40 or 50 years. But someone had to take the plunge.”

Wine country

In January, on Saturday the 24th, around 20 scientists, ethicists, and legal experts traveled to Napa Valley, California, for a retreat among the vineyards at the Carneros Inn. They had been convened by Doudna, the Berkeley scientist who co-discovered the CRISPR system a little over two years ago. She had become aware that scientists might be thinking of crossing the germ line, and she was concerned. Now she wanted to know: could they be stopped?

“We as scientists have come to appreciate that CRISPR is incredibly powerful. But that swings both ways. We need to make sure that it’s applied carefully,” Doudna told me. “The issue is especially human germ-line editing and the appreciation that this is now a capability in everyone’s hands.”

At the meeting, along with ethicists like Greely, was Paul Berg, a Stanford biochemist and Nobel Prize winner known for having organized the Asilomar Conference, a historic 1975 forum at which biologists reached an agreement on how to safely proceed with recombinant DNA, the newly discovered method of splicing DNA into bacteria.

Should there be an Asilomar for germ-line engineering? Doudna thinks so, but the prospects for consensus seem dim. Biotechnology research is now global, involving hundreds of thousands of people. There’s no single authority that speaks for science, and no easy way to put the genie back in the bottle. Doudna told me she hoped that if American scientists agreed to a moratorium on human germ-line engineering, it might influence researchers elsewhere in the world to cease their work.

Doudna said she felt that a self-imposed pause should apply not only to making gene-edited babies but also to using CRISPR to alter human embryos, eggs, or sperm—as researchers at Harvard, Northeastern, and OvaScience are doing. “I don’t feel that those experiments are appropriate to do right now in human cells that could turn into a person,” she told me. “I feel that the research that needs to be done right now is to understand safety, efficacy, and delivery. And I think those experiments can be done in nonhuman systems. I would like to see a lot more work done before it’s done for germ-line editing. I would favor a very cautious approach.”

Not everyone agrees that germ-line engineering is such a big worry, or that experiments should be padlocked. Greely notes that in the United States, there are piles of regulations to keep lab science from morphing into a genetically modified baby anytime soon. “I would not want to use safety as an excuse for a non-safety-based ban,” says Greely, who says he pushed back against talk of a moratorium. But he also says he agreed to sign Doudna’s letter, which now reflects the consensus of the group. “Although I don’t view this as a crisis moment, I think it’s probably about time for us to have this discussion,” he says.

(After this article was published online in March, Doudna’s editorial appeared in *Science*. Along with Greely, Berg, and 15 others, she called for a global moratorium on any effort to use CRISPR to generate gene-edited children until researchers could determine “what clinical applications, if any, might in the future be deemed permissible.” The group, however, endorsed basic research, including applying CRISPR to embryos. The final list of signatories included Church, although he did not attend the Napa meeting.)

As news has spread of germ-line experiments, some biotechnology companies now working on CRISPR have realized that they will have to take a stand. Nessan Bermingham is CEO of Intellia Therapeutics, a Boston startup that raised $15 million last year to develop CRISPR into gene therapy treatments for adults or children. He says germ-line engineering “is not on our commercial radar,” and he suggests that his company could use its patents to prevent anyone from commercializing it.

“The technology is in its infancy,” he says. “It is not appropriate for people to even be contemplating germ-line applications.”

Bermingham told me he never imagined he’d have to be taking a position on genetically modified babies so soon. Modifying human heredity has always been a theoretical possibility. Suddenly it’s a real one. But wasn’t the point always to understand and control our own biology—to become masters over the processes that created us?

Doudna says she is also thinking about these issues. “It cuts to the core of who we are as people, and it makes you ask if humans should be exercising that kind of power,” she told me. “There are moral and ethical issues, but one of the profound questions is just the appreciation that if germ-line editing is conducted in humans, that is changing human evolution.” One reason she feels the research should slow down is to give scientists a chance to spend more time explaining what their next steps could be.

“Most of the public,” she says, “does not appreciate what is coming.”

Antonio Regalado is MIT Technology Review*’s biomedicine editor.*

This barge sucks gas from the lake and sends it to a power plant on shore. It is a small-scale pilot version of the extraction facility in the works.

Lake Kivu's Great Gas Gamble

In a first-of-its-kind endeavor, electricity-starved Rwanda and the Democratic Republic of Congo are trying to get power from a lake—and avert catastrophe.

By Jonathan W. Rosen

Photographs by Jason Florio

It's a Friday afternoon on the Rwandan side of Lake Kivu, and in what was once a quiet cove, a daring venture is taking shape.

Floating just offshore, like a giant mechanical swan, is a nearly completed gas extraction platform: 3,000 tons of concrete and stainless steel that will soon begin capturing a resource not found at this scale in any other lake in the world. Dissolved within Kivu, which straddles the border of Rwanda and the Democratic Republic of Congo (DRC), are approximately 60 billion cubic meters of methane and 300 billion cubic meters of carbon dioxide. The gases, which come from nearby volcanic activity and bacteria decomposing organic material in the lake, represent both danger and economic potential.

If extracted, Kivu's methane could be used to add up to 960 megawatts of electricity-generating capacity, more than six times what Rwanda has now. For both Rwanda and the eastern DRC, which face crippling power shortages and limited options for expanding their electric grids, that could be an economic game changer, supporting new industries and offering a chance to alleviate searing poverty. If the extraction is done properly and the countries can coöperate, it could even help improve their troubled relations and advance stability in a region long beset by turmoil.

Just as critical, removing Kivu's methane may prevent a possible catastrophe. With methane concentrations rising, scientists warn that Kivu will eventually experience a deadly phenomenon known as an overturn. Also known as a limnic eruption, an overturn can occur if the pressure of the gases in a lake exceeds the pressure of the water at a given depth, causing a chain reaction that releases them with violent results. Only two limnic eruptions are known to have occurred in recorded history—both in small lakes in Cameroon in the 1980s. In the deadlier of the two episodes, at Lake Nyos in 1986, more than 1,700 people were asphyxiated when a cloud of carbon dioxide, which burst from the lake along with a 100-meter fountain of water, spread as far as 25 kilometers from shore. Kivu contains a thousand times more gas than Nyos: if even part of it escaped this way, more than two million people living near its shores would be at risk.

In Kivu it's the methane, rather than the carbon dioxide, that's most likely to trigger a gas eruption. That adds urgency to the prospect of harnessing its energy potential, something both Rwanda and the DRC have long sought to do. After decades of little or no progress, gas extraction efforts in both countries have finally gained momentum. On my visit to the lake in February, more than a hundred orange-vested workers were putting the final touches on the first phase of KivuWatt, a $200 million project owned by the U.S. energy firm Contour Global. The lake's first industrial-scale gas-fueled power project, it is expected to add 25 megawatts of generating capacity by the middle of this year and eventually scale up to 100. Another U.S. company, Symbion Power, is set to begin construction of a 50-megawatt project on the Rwandan side of the lake by the end of the year. In the DRC's distant capital, Kinshasa, the Ministry of Hydrocarbons is now reviewing bids for that country's first Kivu gas concession.

Getting the gas out correctly, however, will be tricky. Although the Rwandan government has operated a pilot gas-fired power plant at the lake since 2008, the process of extraction is novel and has been done only on a very small scale. While most experts agree that the lake's methane should be kept from accumulating further in order to prevent a disaster by the end of the century, a few warn that certain extraction processes could disturb the natural stratification that keeps the bulk of the gases trapped in deep waters. Undertaking them could increase, rather than mitigate, the risk of gas eruption. Until a large-scale extraction operation has commenced, it also remains unclear how efficiently the technology will function and how much electricity Kivu will ultimately yield.

"We are very curious to see how our process works," says Jarmo Gummerus, a Finnish engineer and KivuWatt's Rwanda country manager. "Very soon we'll have a much better idea of the potential of this lake."

Boosting the grid

Three hours by car over winding roads from KivuWatt, the Rwandan capital, Kigali, does not appear to be a city in the midst of an energy crisis. In the 21 years since the Rwandan genocide, in which an estimated 800,000 people were killed, the city of a million has transformed from a corpse-ridden backwater into a tidy modern metropolis. Today, Kigali is a town of smooth tree-lined streets, sprouting office towers and American-style subdivisions that stretch to the surrounding hills. It's also the engine of a Rwandan economy that's grown at an average of 8 percent per year over the last decade—one of the highest rates in the world.

As Rwanda and its capital have developed, however, the country's electricity grid has struggled to keep pace. Although installed capacity has doubled in the last five years, it remains a scant 156 megawatts. Today, nearly 80 percent of Rwanda's 12 million people, including the vast majority of rural residents, still lack a connection to the grid. Families and business that do have power, meanwhile, face some of the highest electricity prices in the region—in part because nearly a third of the country's power is generated from imported diesel and

A safety inspector examines the new barge on shore in Kibuye, Rwanda.

Kivu, 1,460 meters above sea level, is part of a system of lakes along the Great Rift Valley.

heavy fuel oil, which arrive by truck from Kenya and Tanzania. According to the World Bank, Rwandan companies pay an average of 24 cents per kilowatt-hour, compared with 15 cents in Kenya and 17 cents in Uganda. The average industrial user in the United States pays less than seven cents.

Hoping to reduce its widespread poverty and boost its small industrial base, Rwanda has set ambitious electrification targets. The country's second Economic Development and Poverty Reduction Strategy, launched in 2013, assumed a nearly fourfold expansion of the power grid, to 563 megawatts, by the end of 2017. Given financial constraints and limited domestic energy resources, however, this will be difficult to pull off. Aside from KivuWatt, the only significant power project nearing completion is a 15-megawatt plant that will burn peat. Although work has begun on another 80-megawatt peat facility, and financing is being arranged for two large-scale regional hydroelectric projects, it's not clear if any will be on the grid by the 2017 target. Rwanda might also have significant geothermal resources, if preliminary surveys are correct, but two exploratory wells drilled in 2013 came up empty. And although Rwanda recently inaugurated East Africa's first utility-scale solar field and authorities are working to bring off-grid solar installations to rural homes, schools, and hospitals, it's unlikely that solar will be able to meet a significant portion of industry's demands. Out of desperation, Rwanda could soon become a significant electricity importer. According to the Ministry of Infrastructure, arrangements are in the works to purchase 30 megawatts from Kenya this year and, eventually, up to 400 megawatts from Ethiopia.

Across the border in the eastern part of the Democratic Republic of Congo (formerly known as Zaire), the power crisis is even more acute. The DRC, a country of 77 million people in a territory roughly the size of Western Europe, contains extensive hydroelectric resources. If fully tapped, the Congo River's Inga Falls could yield an estimated 40,000 megawatts, nearly twice the capacity of the world's largest power station, the Three Gorges Dam in China. Today, however, the DRC's aging grid has an installed capacity of just 2,400 megawatts, roughly half of which is routinely unavailable because the transmission infrastructure is in such poor shape. In the war-torn east, power is particularly limited. Goma, the largest city on Lake Kivu, has an available capacity of less than five megawatts—a meager amount for a town of a million residents and a situation, some argue, that helps promote conflict. If boosting eastern Congo's grid can spur the development of industries, says Bantu Lukambo, an environmental activist based in Goma, that would reduce the appeal of the region's dozens of armed groups, which are magnets for youth with no other employ-

Gases from the lake will enter the gray stainless-steel tube to be separated.

ment prospects. In addition, he says, more development could weaken the market for illicit charcoal, a trade that generates millions of dollars per year for local militias and leads to extensive deforestation.

The volcanoes responsible for much of Kivu's gas loom over Goma and its environs. In 2002, an eruption of Nyiragongo, a volcano located 20 kilometers north of town, destroyed a fifth of the city, leaving tens of thousands of people homeless and depositing lava that's still being used as a building material. On a drive west from town, Mathieu Yalire, chief geochemist at the government-run Goma Volcano Observatory, shows me several depressions known to contain *mazuku*, lethal seepages of carbon dioxide that are concentrated near the ground at the edges of past lava flows and occasionally asphyxiate children. At Kasinga Primary School in the town of Sake, 25 kilometers west of Goma, principal Batchoka Lubungo shows us a photo, displayed on the wall of his office, of a young *mazuku* victim.

"One morning we found the boy dead over there," he says, pointing to a known danger zone just outside his window. "We keep this picture here as a warning to the students."

The presence of *mazuku* is a reminder of Lake Kivu's potential to sow disaster. But the carbon dioxide is not the only danger. The lake's geochemistry is unusual, largely as a consequence of local subaquatic springs that absorb carbon dioxide from the region's volcanic soil and feed the gas into Kivu's deepest waters. Much of the methane comes from decomposing organic matter; the rest comes either from the volcanic

soil or from bacteria converting the carbon dioxide to methane. Critically, these springs are saline, while the water sources feeding the lake's upper layers are fresh. Since saline water is much denser than fresh water, this creates density gradients that prevent the gases from diffusing upward and into the atmosphere. Although this stratification is stable at present, the gas accumulation it makes possible has apparently led to limnic eruptions in the distant past. If nothing is done, it is likely to do so in the future.

Still, much about this risk remains uncertain. Studies of Kivu's sediment record suggest that the lake has experienced at least five overturns in the last 6,000 years. It's not clear, however, whether these events involved all the lake's layers of water, thus releasing all its gas, or just portions of its upper layers. In addition, though recent measurements have found that the concentration of methane is increasing—at a rate that could bring the gas close to saturation by the end of the century—it's not yet known why this is happening or whether it will continue. Complicating matters, Kivu consists of five different basins of varying depths, each with distinct physiochemical properties.

It is clear, though, that an eruption in Kivu's main basin could cause a disaster of apocalyptic proportions. If all the methane and carbon dioxide currently dissolved in Kivu were released into the atmosphere, they would cover the entire lake in a cloud of gas more than 100 meters thick. If even a small fraction of the gas were to get out, it could suffocate entire towns along the lake shore. This can happen if water at a given depth becomes fully saturated with gas and is lifted by a big earthquake, a volcanic eruption, or another external disturbance to a depth where the water pressure is not great enough to keep the gas dissolved.

Extraction

Whatever the extent of Kivu's eruption risks, its methane has long been of commercial interest. From 1963 until 2006, Rwanda's lakeside Bralirwa Brewery fueled its boilers with methane extracted 800 meters offshore. In the 1980s, researchers from the Netherlands tapped Bralirwa's excess gas to fuel a fleet of cars, though the project eventually foundered. By the early 1990s, cross-border efforts to utilize the gas for electricity had begun to gain momentum, but progress was cut short by the Rwandan genocide and subsequent wars in eastern Congo. Eventually, with the return of stability in Rwanda, the government in Kigali entered into partnership with a Scottish firm to build a small pilot facility on the lake. The plant, known as KP1, began working intermittently in 2008 and produces a few megawatts of electricity. Another project briefly produced 2.4 megawatts in 2010, but the equipment was removed from the lake after it was damaged in a storm.

KivuWatt is the country's first attempt at large-scale gas extraction. Even though its technology is novel, the concept is relatively simple. A barge will be anchored to the lake bed 13 kilometers offshore, where four plastic pipes will draw up water from 350 meters below the surface. As the water rises, bubbles of methane and carbon dioxide will begin to form; eventually, roughly 80 percent of the methane and 40 percent of the carbon dioxide will be siphoned off inside a subsurface horizontal chamber known as a separator. From there, the partially degassed water will be reinjected deep into the lake, and the gas—at this point roughly 30 percent methane and 70 percent carbon dioxide, with trace amounts of hydrogen sulfide—will continue upward into one of four towers on the barge. Here, "wash water" taken from a depth of 40 meters will be mixed with the gas to remove as much of the remaining carbon dioxide as possible. This water will be returned to a depth of 60 meters, shallow enough for some of the carbon dioxide to eventually diffuse into the atmosphere. The end product, a gas composed of roughly 85 percent methane, will then be pressurized and sent to a power plant on shore.

For a region in such dire need of electricity, Kivu's gas is an attractive proposition. According to Gummerus, the engineer overseeing the project, KivuWatt will sell power generated in its first phase to Rwanda's state-owned utility for less than 15 cents per kilowatt-hour. That's competitive with the rates expected from the country's forthcoming peat projects and less than half the cost of power generated from imported fossil fuels. (Later the rate is expected to fall to less than 12 cents per kilowatt-hour.)

As desirable as the project seems for both economic and safety reasons, however, it could pose environmental risks of its own—including the chance that the degassing operations could change the structure and properties of the lake.

According to the Management Prescriptions for the Development of Lake Kivu Gas Resources, a 2009 document—known as the MPs—that both Rwanda and the DRC have adopted as general guidelines for gas extraction, the risks have less to do with the removal of the gas itself than with the reinjection of the degassed water. Since Kivu's methane-rich deep water is saline, dense, and abundant in nutrients, releasing it near the surface could damage the lake's ecosystem and weaken its density stratification. To mitigate these risks, the report dictates that all water extracted from the lake's deep-water "resource zone" be reinjected at least 260 meters below the surface so that it remains under enough pressure. Because KivuWatt's design was approved before the report

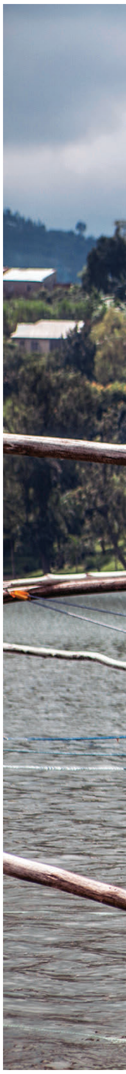

Top, women dry tiny sambaza*, or sardines. Bottom, a sign in Kibuye explains KivuWatt.*

Children of fishermen play around fishing boats off the shores of Gisenyi.

Workers attend to the final construction phase of the KivuWatt barge.

was released, however, the project isn't subject to this requirement, and its degassed water will be returned above that level, at a depth of 240 meters. Although this distinction may seem trivial, Philip Morkel, a South African engineer and a member of the five-person expert committee that wrote the guidelines, believes otherwise. "Once you punch through those gradient layers, you start damaging the protective mechanism that the lake has to preserve itself," he says. "On a large scale it becomes seriously problematic."

Some experts see less reason for alarm. Dario Tedesco, an Italian volcanologist with extensive knowledge of the lake, tells me the quantity of water reinjected is unlikely to be large enough to create a serious disturbance. Alfred Johny Wüest, head of the aquatic physics research group at the Swiss Federal Institute of Aquatic Science and Technology and another member of the MPs committee, says it erred on the side of caution, which means that even the water from KivuWatt will be reinjected deep enough for safety.

Wüest has other concerns, however, including that the project could harm the lake's ecology. That worries some fishermen, too. Fishing will be forbidden in an exclusion zone around KivuWatt, but even outside that area, the reinjection of wash water and the noise and vibrations from the extraction project could be noticeable. Fishermen near the KP1 pilot plant told me that their catches, mainly of sardines known as *sambaza*, have fallen significantly since the facility began to operate. It's difficult to isolate the cause there, however, since a drop in *sambaza* numbers has been documented across the

A portrait of Rwanda's president looms over controls in the KP1 pilot plant.

lake—a decline linked to a rise in unregulated boats and the introduction of a predator species.

The most pressing concern from an energy standpoint is how well the gas extraction technology will perform. The MPs' authors estimate that the lake is capable of providing between 160 and 960 megawatts of generating capacity over a period of 50 years. After that, the gas that continues to accumulate could be harnessed for up to 100 megawatts of capacity in perpetuity.

All of this, though, hinges on the efficiency of the extraction process and the onshore power conversion technology. KivuWatt should extract about two-thirds of the methane from the water it draws up, with the rest lost during the separation process and in the washing towers. Although some of that gas will be returned to the lake in the reinjected water and could theoretically be reëxtracted at a later date, it is likely to settle in smaller concentrations, which could make it uneconomical to capture. But these projections are based only on simulations, and the real efficiency won't be known until the barge begins to operate. Essentially, the lake's potential is still highly uncertain.

Disappointing results could spur disputes between Rwanda and the Democratic Republic of Congo. Although the two countries signed an agreement in 1975 to share the methane equally, Rwanda's head start in exploitation means it could potentially infringe upon its neighbor's resource, particularly if the lake yields less power than expected. (Because gas levels in the lake's main basin are uniform across a given depth, it's

impossible to extract from Rwandan waters without affecting the quantity available in the DRC.) The two countries do not exactly have an amicable history. Since the current Rwandan government seized power at the end of the genocide, Rwanda has twice launched rebellions to topple governments in Kinshasa—once successfully. It has also been implicated in the support of numerous proxy militias in the DRC's east as well as the widespread smuggling of Congolese minerals, including gold, tin, and coltan, an ore mined for use in electronics manufacturing. In this context, some worry, Kivu's gas may simply be another means by which Rwanda, a small yet highly organized state, manages to profit at the expense of its larger, dysfunctional neighbor.

If the project is successful, though, it could help mend cross-border relations. The two countries, which already share power from a hydroelectric plant on their border south of Kivu, have long sought to implement a joint gas-to-power project, according to Augusta Umutoni, head of Rwanda's Lake Kivu Monitoring Program, a government body that oversees the extraction process. Despite the vagaries of politics, she says, energy officials from the two countries collaborate well at a technical level. Both, however, are waiting until the concept has been shown to work on a commercial scale.

With so many questions lingering as KivuWatt prepares for launch, it's easy to forget that KP1, the pilot facility, is now in its eighth year of operation. Although the project, originally conceived to generate five megawatts, struggled to produce more than a single megawatt for years, it is now consistently generating between two and three—giving authorities a slight boost of confidence.

On a February evening, I join Olivier Ntirushwa, KP1's plant manager, on a tour of the barge, which floats a kilometer from shore, within sight of the Rwanda-DRC border. Looking north, I see the city of Goma and its erratic patchwork of lights. Further afield is the steaming cone of Nyiragongo. Over the noise of the barge's machinery, Ntirushwa gives me a brief history of the project: how it operated well below capacity for years; how recent improvements to its separator finally boosted its power output. When I ask about his predictions for KivuWatt, Ntirushwa says he doesn't know much about the project, or how much electricity the lake might eventually yield. Instead, he stresses how it feels to be involved in such a high-stakes experiment.

"It's exciting because we are the pioneers of this technology," he says as we stare out over the water. "Nobody else has ever done this."

Jonathan W. Rosen is a journalist based in Kigali.

After the barge is completed, it will be towed 13 kilometers onto the lake to begin extracting gas.

LEO ESPINOSA

Machine Dreams

To rescue its struggling business, Hewlett-Packard is making a long-shot bid to change the fundamentals of how computers work.

By Tom Simonite

There is a shrine inside Hewlett-Packard's headquarters in Palo Alto, in the heart of Silicon Valley. At one edge of HP's research building, two interconnected rooms with worn midcentury furniture, vacant for decades, are carefully preserved. From these offices, William Hewlett and David Packard led HP's engineers to invent breakthrough products, like the 40-pound, typewriter-size programmable calculator launched in 1968.

In today's era of smartphones and cloud computing, HP's core products could also look antiquated before long. Revenue and profit have slid significantly in recent years, pitching the company into crisis. HP is sustained mostly by sales of servers, printers, and ink (its PCs and laptops contribute less than one-fifth of total profits). But businesses have less need for servers now that they can turn to cloud services run by companies like Amazon—which buy their hardware from cheaper suppliers than HP. Consumers and businesses rely much less on printers than they once did and don't expect to pay much for them.

HP has shed over 40,000 jobs since 2012, and it will split into two smaller but similarly troubled companies later this year (an operation that will itself cost almost $2 billion). HP Inc. will sell printers and PCs; Hewlett Packard Enterprise will offer servers and information technology services to corporations. That latter company will depend largely on a division whose annual revenue dropped by more than 6 percent between 2012 and 2014. Earnings shrank even faster, by over 20 percent. IBM, HP's closest rival, sold off its server business to China's Lenovo last year under similar pressures.

And yet, in the midst of this potentially existential crisis, HP Enterprise is working on a risky research project in hopes of driving a remarkable comeback. Nearly three-quarters of the people in HP's research division are now dedicated to a

single project: a powerful new kind of computer known as "the Machine." It would fundamentally redesign the way computers function, making them simpler and more powerful. If it works, the project could dramatically upgrade everything from servers to smartphones—and save HP itself.

"People are going to be able to solve problems they can't solve today," says Martin Fink, HP's chief technology officer and the instigator of the project. The Machine would give companies the power to tackle data sets many times larger and more complex than those they can handle today, he says, and perform existing analyses perhaps hundreds of times faster. That could lead to leaps forward in all kinds of areas where analyzing information is important, such as genomic medicine, where faster gene-sequencing machines are producing a glut of new data. The Machine will require far less electricity than existing computers, says Fink, making it possible to slash the large energy bills run up by the warehouses of computers behind Internet services. HP's new model for computing is also intended to apply to smaller gadgets, letting laptops and phones last much longer on a single charge.

It would be surprising for any company to reinvent the basic design of computers, but especially for HP to do it. It cut research jobs as part of downsizing efforts a decade ago and spends much less on research and development than its competitors: \$3.4 billion in 2014, 3 percent of revenue. In comparison, IBM spent \$5.4 billion—6 percent of revenue—and has a much longer tradition of the kind of basic research in physics and computer science that creating the new type of computer will require. For Fink's Machine dream to be fully realized, HP's engineers need to create systems of lasers that fit inside fingertip-size computer chips, invent a new kind of operating system, and perfect an electronic device for storing data that has never before been used in computers.

Pulling it off would be a virtuoso feat of both computer and corporate engineering.

New blueprint

In 2010, Fink was running the HP division that sells high-powered corporate servers and feeling a little paranoid. Customers were flocking to startups offering data storage based on fast, energy-efficient flash memory chips like the ones inside smartphones. But HP was selling only the slower, established storage technology of hard drives. "We weren't responding aggressively enough," says Fink. "I was frustrated we were not thinking far enough into the future."

Trying to see a way to leapfrog ahead, he wondered: why not use new forms of memory not just to upgrade data storage but to reinvent computers entirely? Fink knew that researchers at HP and elsewhere were working on new memory technologies that promised to be much faster than flash chips. He and his chief technical advisor drew up a blueprint that would use those technologies to make computers far more powerful and energy efficient.

An internal paper on that idea, with the corporatese title "Unbound Convergence," went nowhere. But when Fink was appointed HP's chief technology officer and head of HP Labs two years later, he saw a chance to resurrect his proposal. "When I looked at all the teams in Labs, I could see the pieces were there," he says. In particular, HP was working on a competitor to flash, based on a device called the memristor. Although the memristor was still in development, it looked to Fink as if it would someday be fast enough and store data densely enough to make his blueprint realizable. He updated his previous proposal and gave the computer a name he thought would be temporary: the Machine. It stuck.

The Machine is an attempt to update the design that has defined the guts of computers since the 1970s. Essentially, computers are constantly shuttling data back and forth between different pieces of hardware that hold information.

The HP Way

1
1938: Hewlett and Packard set up shop in a rented garage. Their first product is an audio oscillator used to test sound equipment.

2
1951: Devices that engineers use to measure the frequency of electronic signals establish HP's second major business and help the expanding nuclear industry.

1964: HP launches an atomic clock that accurately counts time by millionths of a second; it is used to synchronize international time standards.

1


2


COURTESY OF HP

1966: HP's first computer is created to provide a way to process data from scientific instruments. The cabinet-size machine starts at $25,000.

1972: The first scientific calculator small enough to fit into a shirt pocket is launched and becomes a hit product that helps kill off the slide rule.

1984: The world's first desktop laser printer, the HP LaserJet, is launched and quickly becomes a staple of office life.

1995: Chemist Stanley Williams establishes HP's first fundamental research program, investigating whether organic molecules could be smaller replacements for silicon transistors.

2008: Williams's group creates a tiny electronic device called the memristor, inventing a promising new form of data storage.

2012: New CTO Martin Fink puts most HP researchers to work on a new design for computers based on memristor memory. The first "Machine" should be complete sometime before 2020.

3


4


5


6


PHOTO OF MARTIN FINK COURTESY OF RICHARD LEWINGTON, HP LABS

One, known as storage, keeps your photos and documents plus the computer's operating system. It consists of hard drives or flash memory chips, which can fit a lot of data into a small space and retain it without power (engineers call it "nonvolatile" memory). But both hard drives and flash chips read and write data very slowly relative to the pace that a computer's processor can work on it. When a computer needs to get something done, the data must be copied into short-term memory, which uses a technology 10,000 or more times faster: DRAM (dynamic random-access memory). This type of memory can't store data very densely and goes blank when powered down.

The two-tier system of storage and memory means computers spend a lot of time and energy moving data back and forth just to get into a position to use it. This is why your laptop can't boot up instantly: the operating system must be retrieved from storage and loaded into memory. One constraint on the battery life of your smartphone is its need to spend energy keeping data alive in DRAM even when it is idling in your pocket.

That may be a mere annoyance for you, but it's a costly headache for people working on computers that do the sort of powerful number-crunching that's becoming so important in all kinds of industries, says Yuanyuan Zhou, a professor at the University of California, San Diego, who researches storage technologies. "People working on data-intensive problems are limited by the traditional architecture," she says. HP estimates that around a third of the code in a typical piece of data analysis software is dedicated purely to juggling storage and memory, not to the task at hand. This doesn't just curtail performance. Having to transfer data between memory and storage also consumes significant energy, which is a major concern for companies running vast collections of servers, says Zhou. Companies such as Facebook spend millions trying to cut the huge power bills for their warehouses of computers.

The Machine is designed to overcome these problems by scrapping the distinction between storage and memory. A single large store of memory based on HP's memristors will both hold data and make it available for the processor. Combining memory and storage isn't a new idea, but there hasn't yet been a nonvolatile memory technology fast enough to make it practical, says Tsu-Jae King Liu, a professor who studies microelectronics at the University of California, Berkeley. Liu is an

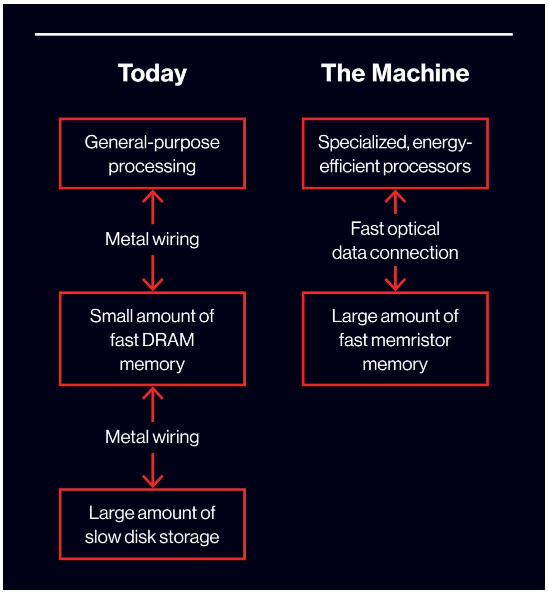

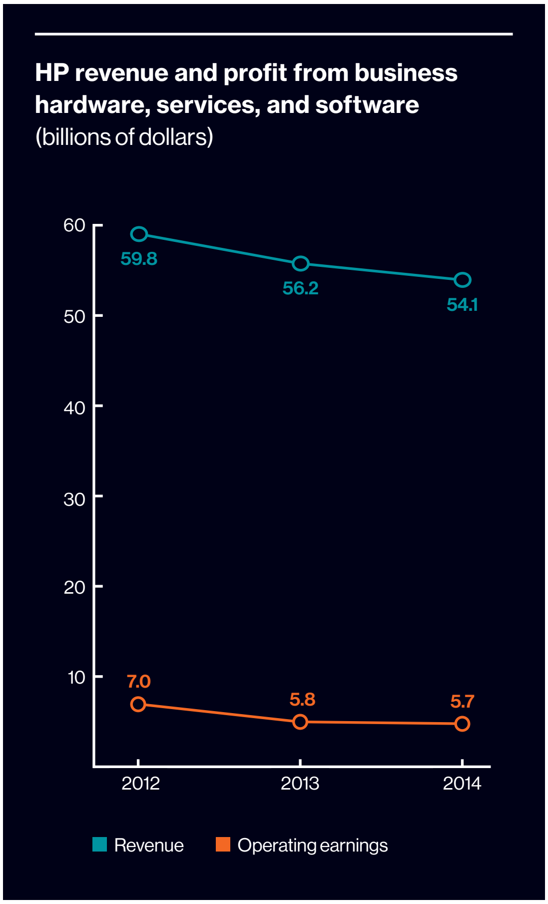

advisor to Crossbar, a startup working on a memristor-like memory technology known as resistive RAM. It and a handful of other companies are developing the technology as a direct replacement for flash memory in existing computer designs. HP is alone, however, in saying its devices are ready to change computers more radically.

To make the Machine work as well as Fink imagines, HP needs to create memristor memory chips and a new kind of operating system designed to use a single, giant store of memory. Fink's blueprint also calls for two other departures from the usual computer design. One is to move data between the Machine's processors and memory using light pulses sent over optical fibers, a faster and more energy-efficient alternative to metal wiring. The second is to use groups of specialized energy-efficient chips, such as those found in mobile devices, instead of individual, general-purpose processors. The low-energy processors, made by companies such as Intel, can be bought off the shelf today. HP must invent everything else.

Primary occupation

No one has built a fundamentally new operating system for decades. For more than 40 years, every "new" operating system has actually been built or modeled on one that came before, says Rich Friedrich, director of system software for the Machine. Academic research on operating systems has been minimal because existing systems are so dominant.

The work of Friedrich and his colleagues will be crucial. The software must draw together the Machine's various unusual components into a reliable system unlike any other computer ever built. The group also has to help market that computer. If the operating system doesn't look attractive to other companies and programmers, the Machine's technical merits will be irrelevant. For that reason, HP is working on two new operating systems at the same time. One is based on the widely used Linux system and will be released this summer, along with software to emulate the hardware it needs to run. Linux++, as it is called, can't make full use of the Machine's power but will be compatible with most existing Linux software, so programmers can easily try it out. Those who like it will be able to step up to HP's second new operating system, Carbon, which won't be finished for two years or more. It will be released as open source, so anyone can inspect or modify its code, and is being designed from the ground up to unleash the full power of a computer with no division between storage and memory. By starting from scratch, Friedrich says, this operating system will remove all the complexity, caused by years of updates on top of updates, that leads to crashes and security weaknesses.

Packard and Hewlett in 1964.

If everything goes to plan, Hewlett Packard Enterprise will be five years old when the first version of the Machine can ride to its rescue.

COURTESY OF HP

Tests with the closest thing to a working version of the Machine—a simulation running inside a cluster of powerful servers—hint at what Carbon and the new computer might be able to deliver once up and running. In one trial, the simulated Machine and a conventional computer raced to analyze a photo and search a database of 80 million other images to find the five that were most visually similar. The off-the-shelf, high-powered HP server completed the task in about two seconds. The simulated Machine needed only 50 milliseconds.

Handling such tasks tens or hundreds of times faster—with the same energy expenditure—could be a crucial advantage at a time when more and more computing problems involve huge data sets. If you have your genome sequenced, it takes hours for a powerful computer to refine the raw data into the finished sequence that can be used to scrutinize your DNA. If the Machine shortened that process to minutes, genomic research could move faster, and sequencing might be easier to use in medical practice. Sharad Singhal, who leads HP's data analysis research, expects particularly striking improvements for problems involving data sets in the form of a mathematical graph—where entities are linked by a web of connections, not organized in rows and columns. Examples include the connections between people on Facebook or between the people, planes, and bags being moved around by an airline. And wild new applications are likely to emerge once the Machine is working for real. "Techniques that we would discard as impractical suddenly become practical," Singhal says. "People will think of ways to use this technology that we cannot think of today."

Perfecting the memristor is crucial if HP is to deliver on that striking potential. That work is centered in a small lab, one floor below the offices of HP's founders, where Stanley Williams made a breakthrough about a decade ago.

Williams had joined HP in 1995 after David Packard decided the company should do more basic research. He came to focus on trying to use organic molecules to make smaller, cheaper replacements for silicon transistors (see "Computing After Silicon," September/October 1999). After a few years, he could make devices with the right kind of switchlike behavior by sandwiching molecules called rotaxanes between platinum electrodes. But their performance was maddeningly erratic. It took years more work before Williams realized that the molecules were actually irrelevant and that he had stumbled into a major discovery. The switching effect came from a layer of titanium, used like glue to stick the rotaxane layer to the electrodes. More surprising, versions of the devices built around that material fulfilled a prediction made in 1971 of a completely new kind of basic electronic device. When Leon Chua, a professor at the University of California,

Berkeley, predicted the existence of this device, engineering orthodoxy held that all electronic circuits had to be built from just three basic elements: capacitors, resistors, and inductors. Chua calculated that there should be a fourth; it was he who named it the memristor, or resistor with memory. The device's essential property is that its electrical resistance—a measure of how much it inhibits the flow of electrons—can be altered by applying a voltage. That resistance, a kind of memory of the voltage the device experienced in the past, can be used to encode data.

HP's latest manifestation of the component is simple: just a stack of thin films of titanium dioxide a few nanometers thick, sandwiched between two electrodes. Some of the layers in the stack conduct electricity; others are insulators because they are depleted of oxygen atoms, giving the device as a whole high electrical resistance. Applying the right amount of voltage pushes oxygen atoms from a conducting layer into an insulating one, permitting current to pass more easily. Research scientist Jean Paul Strachan demonstrates this by using his mouse to click a button marked "1" on his computer screen. That causes a narrow stream of oxygen atoms to flow briefly inside one layer of titanium dioxide in a memristor on a nearby silicon wafer. "We just created a bridge that electrons can travel through," says Strachan. Numbers on his screen indicate that the electrical resistance of the device has dropped by a factor of a thousand. When he clicks a button marked "0," the oxygen atoms retreat and the device's resistance soars back up again. The resistance can be switched like that in just picoseconds, about a thousand times faster than the basic elements of DRAM and using a fraction of the energy. And crucially, the resistance remains fixed even after the voltage is turned off.

When Williams announced the memristor in 2008, it began what he now calls a roller coaster ride, because the basic research finding quickly progressed to become a crucial development project for HP. "Sometimes the adrenaline can be a little overwhelming," he says. Memristors have been his team's primary occupation since 2008, and in 2010 HP announced that it had struck a deal with the South Korean memory chip manufacturer SK Hynix to commercialize the technology. At the time, HP was focused on getting memristors to replace flash memory in conventional computers. Then, in 2012, Fink piled on more pressure by putting memristors at the heart of his blueprint for the Machine.

Risky business

Using memristors, whether as a replacement for flash memory or as the basis for the Machine, would require packaging them into memory chips that combine a dense array of the

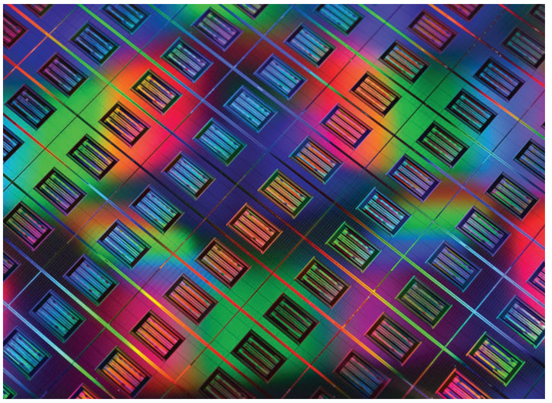

Top, a close-up of memristors on a silicon wafer. Below, a detail in a prototype of the Machine.

Stanley Williams says the memristor will offer an unrivaled combination of speed, density, and energy efficiency.

COURTESY OF HP

devices with conventional silicon control circuitry. But such chips don't exist at the moment. And it's not clear when HP will be able to get hold of any. It is up to a chip maker such as SK Hynix to develop reliable memristor chips that are suited to its production lines. So far it sends HP silicon wafers with memristors that can be tested individually, not used in a computer.

Fink and Williams say that the first prototype memory chips could arrive sometime next year. (Previous statements by Williams made it seem as if the technology would hit the market in 2013, but he and Fink say these remarks were misinterpreted.) A spokesman for Hynix, Heeyoung Son, declined to comment on whether 2016 is feasible. "There is no specific time line," he said. "It will take some time to reach that."

Edwin Kan, a professor at Cornell University who works on memory technology, says that progress on memristors and similar devices appeared to stall when companies tried to integrate them into dense, reliable chips. "It looks promising, but it has been looking promising for too long," he says.

Dmitri Strukov, one of Williams's former collaborators at HP, says memristors have yet to pass a key test. Strukov, an assistant professor at the University of California, Santa Barbara, and lead author on the 2008 paper announcing the memristor, says that while technical publications released by HP and SK Hynix have shown that individual memristors can be switched trillions of times without failing, it's not yet clear that large arrays perform the same way. "That's nontrivial," he says.

Surrounded by prototype memristors in his lab, Williams cheerily says his confidence in the technology remains high. "If I had thought something else was coming along that was better, I would have jumped on that in a femtosecond," he says. Williams keeps a chart on the wall summarizing the competition from a handful of other memory technologies that companies including IBM and Samsung are positioning as replacements for flash. Those companies and HP are at similar stages: progress is being reported, though the commercial future of their devices is unclear. But Williams says none of the other technologies has as good a combination of speed, density, and energy efficiency as memristors. Even though the first generation of memristor chips has yet to be introduced, his group is already looking at ways for subsequent generations to pack in more and more data. Those include stacking memristors up in layers and storing more than a single bit in each memristor.

In the lab next door, researchers are working on optical data connections compact enough to link up the components inside a computer. But this project is at an even earlier stage. HP's engineers can show off a silicon wafer covered in tiny lasers, each a quarter the width of a human hair, and use them to pulse light down slim optical fibers. The lasers are intended to be packaged into small chips and added to the Machine's circuit boards to connect its memory and processors. Existing computers use metal wires for that, because conventional optical connection technology is too bulky. Fink expects this to be the last of the Machine's components to be ready, around the end of the decade.

When that time comes, the first full version of the Machine could be rented out for companies or universities to access remotely for large-scale data analysis, Fink says. Soon after, Machine-style servers will be available to buy. Work will also begin on scaling down the new computer for use in other types of devices, such as laptops, smartphones, and even cable TV boxes, he says.

If everything goes to plan, Hewlett Packard Enterprise will be five years old when the first version of the Machine can ride to its rescue. But the market may not be much more welcoming than it is to HP's existing business today.

New computing technologies tend to start out expensive and difficult to use. Interest in HP's servers is waning because cloud services offer companies the exact opposite—a cheaper and easier alternative to services that once had to be built in-house. Even companies that buy their own servers are likely to keep seeing significant performance improvements from steady upgrades in technologies such as networking and storage for several years, says Doug Burger, a senior researcher at Microsoft. "A clean sheet is great because you get to redefine everything, but you have to redefine everything," he says. "Disruptive, successful projects in very complex system architectures are pretty rare."

Fink argues that his researchers can create tools that ease the process of migrating to the strange new future of computing. And he thinks companies will become more open to embracing a new platform as incremental upgrades to conventional computers start to yield diminishing returns.

Somewhat paradoxically, Fink also says that the prospect of failing to deliver the Machine at all doesn't trouble him. HP can fall back on selling memristor chips and photonic interconnects as ways to upgrade the flawed architecture of today's computers, he says—though HP isn't in the business of manufacturing components today. "We're gaining a lot by doing this," he says. "It's not all or nothing."

HP's struggles seem to suggest just the opposite.

Tom Simonite is MIT Technology Review*'s San Francisco bureau chief.*

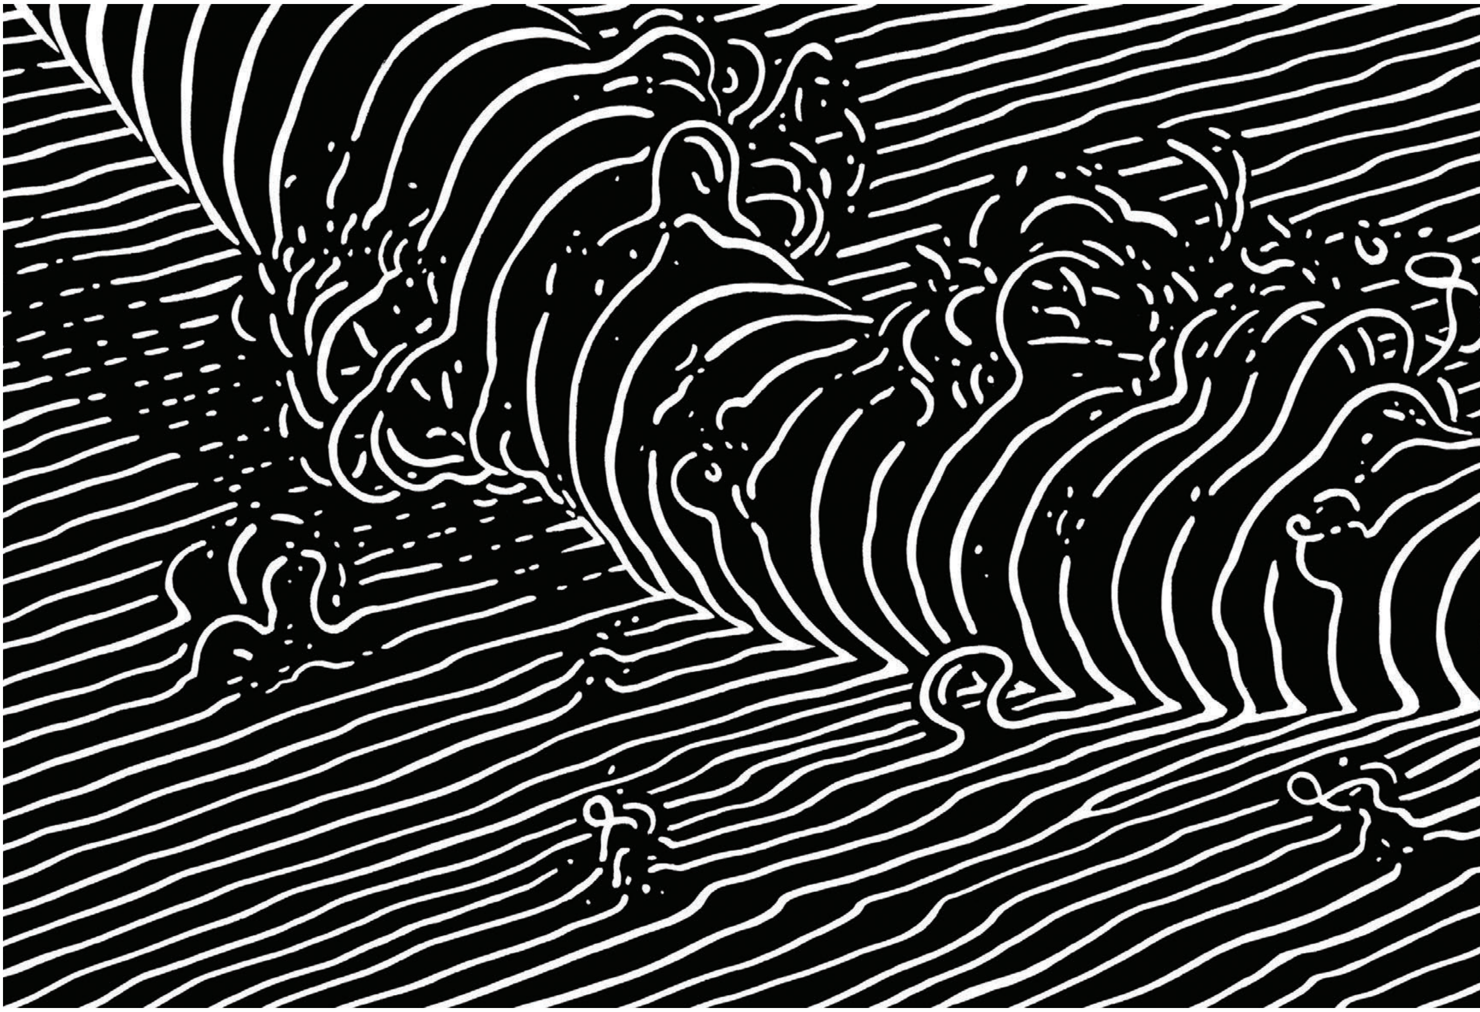

BRENDAN MONROE

Paralyzed Again

We have the technology to dramatically increase the independence of people with spinal-cord injuries. The problem is bringing it to market—and keeping it there.

By Brian Bergstein

One night in 1982, John Mumford was working on an avalanche patrol on an icy Colorado mountain pass when the van carrying him and two other men slid off the road and plunged over a cliff. The other guys were able to walk away, but Mumford had broken his neck. The lower half of his body was paralyzed, and though he could bend his arms at the elbows, he could no longer grasp most things in his hands.

Fifteen years later, however, he received a technological wonder that reactivated his left hand. It was known as the Freehand System. A surgeon placed a sensor on Mumford's right shoulder, implanted a pacemaker-size device known as a stimulator just below the skin on his upper chest, and threaded wires into the muscles of his left arm. On the outside of Mumford's body, a wire ran from the shoulder sensor to an external control unit; another wire ran from that control unit to a transmitting coil over the stimulator in his chest. Out of this kludge came something incredible: by maneuvering his right shoulder in certain ways, Mumford could send signals

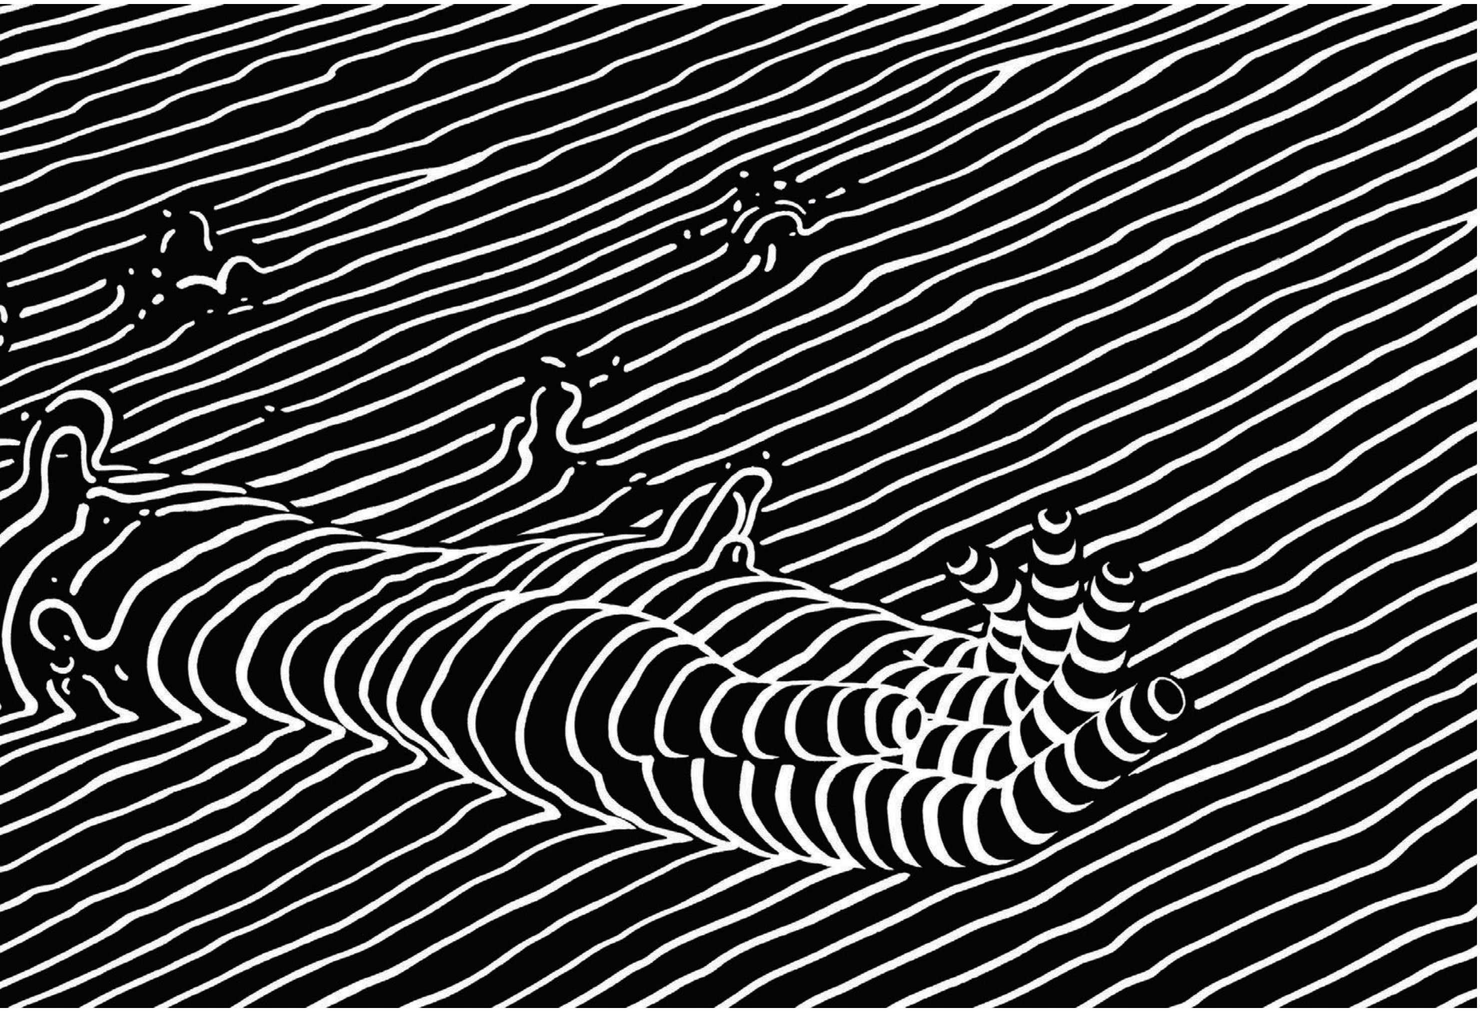

through the stimulator and down his left arm into the muscles of his hand. The device fell short of perfection—he wished he could throw darts with his buddies. But he could hold a key or a fork or a spoon or a glass. He could open the refrigerator, take out a sandwich, and eat it on his own. Mumford was so enthusiastic that he went to work for the manufacturer, a Cleveland-area company called NeuroControl, traveling the country to demonstrate the Freehand at assistive-technology trade shows.

Mumford was in Cleveland for a marketing meeting in 2001 when he got news that still baffles him: NeuroControl was getting out of the Freehand business. It would focus instead on a bigger potential market with a device that helped stroke victims. Before long, NeuroControl went out of business entirely, wiping out at least $26 million in investment. At first, Mumford remained an enthusiastic user of the Freehand, though one thing worried him: the wires running outside his body would sometimes fray or break after catching on clothing. Each time, he found someone who could reach into his supply of replacements and reconnect the system. But by 2010, the last wire was gone, and without the prospect of tech support from NeuroControl, the electrical equipment implanted in Mumford's body went dormant. He lost the independence that had come from having regained extensive use of one hand. "To all of a sudden have that taken away—it's incredibly frustrating," he says. "There's not a day where I don't miss it."

Mumford's voice rises in astonishment as he tells the tale. "I have a device implanted in my body that was considered to be one of the best innovations or inventions of that century," he says. "The last thing you think is that the company is going to go out of business, and not only is it going to go out of business, but you're not even going to be able to buy parts for that. That seems insane!"

Around 250 people are believed to have gotten the Freehand from NeuroControl, and Mumford was far from the only

one heartbroken by the company's failure. Their experience is a cautionary tale now for any implantable medical device that might serve "orphan markets"—relatively small groups of people. Although advances in brain-machine interfaces and electrical-stimulation devices are generating marvelous research results in people with paralysis—some are using their thoughts to control robotic arms, and others are taking tentative steps—it's possible those breakthroughs won't last long on the market, assuming they can be commercialized at all. Limp limbs can be reanimated by technology, but they can be quieted again by basic market economics.

The initial flourish

The technology in Mumford's body began to be developed in the 1970s. The lead inventor, P. Hunter Peckham, a biomedical engineer at Case Western Reserve University in Cleveland, wanted to see whether electrical stimulation would reverse atrophy and ultimately restore function to paralyzed muscles. First in animals and then in people, Peckham and colleagues used hypodermic needles to inject tiny coils of wire into muscles, near nerves. They could then send mild pulses of electricity through these wires and stimulate the muscles, changing their very structure. Over time, by putting the wires in the right places and precisely tuning the bursts of electricity, the researchers could coördinate the muscles' movements—re-creating, among other things, the normal grasp of a hand. Eventually the scientists figured out how to implant the technology into patients and let them operate it themselves, outside the lab, by means of a joystick-like control unit mounted to the shoulder. The first version of what would become the Freehand system was installed in a patient in 1986. Peckham and five other investors founded NeuroControl seven years later with technologies licensed from Case Western.

When the U.S. Food and Drug Administration approved the Freehand in 1997, it was a milestone. It was not the first commercial bionic device—pacemakers and cochlear implants already existed—but it was the first that helped paralyzed patients regain some use of the hands. In fact, it was the first one that used electrical stimulation to make joints move—and to this day it remains the only one ever released.

Independent research showed that even at a cost of around $60,000 (for the device and the necessary surgery), the Freehand saved money in the long run by reducing a patient's need for attendant care. But while the technology was impressive, the Freehand got stuck in a small niche.

Although there are 250,000 people with spinal-cord injuries in the United States alone, the Freehand worked only for people whose paralysis stemmed from an injury to a certain area—between the fifth and sixth vertebrae of their cervical spine. That's because a break in that location left them with enough shoulder and elbow mobility to trigger the Freehand's grasp-and-release function. Although NeuroControl estimated its potential market at more than 50,000 people in the United States, not all of them were willing or healthy enough to endure the major operation that was required to implant the device and all those wires.

Most important, the potential market was further narrowed by the fact that some private insurers and Medicare, the U.S. government insurance program for the elderly and the disabled, would not always cover the full cost. Rehabilitation clinics and hospitals were already likely to be conservative about recommending a novel implantable system to patients. But given that they might absorb any uncovered costs from the procedure, many medical centers were more reluctant to advocate the technology than NeuroControl had hoped.

Lacking momentum, NeuroControl stopped selling the product. "The investors had expected that it would penetrate a much larger volume of the overall spinal-injury population," says Geoff Thrope, who was NeuroControl's director of business development. "We were able to make dozens of implant sales per year. You need to be in the hundreds, if not thousands, to have it make sense."

But the decision still rankles Peckham, who resigned from NeuroControl's board as a result. With some more time, he says, NeuroControl might have seen its way through to a sustainable business. It had 19 patients enrolled in a clinical trial in England; one more would have given it the 20 necessary to allow the British national health-care system to move toward covering the cost of the Freehand. The U.S. Department of Veterans Affairs was likely to follow suit, he says. The problem was that other board members—primarily venture capitalists who "decided they were not seeing the return on the investment they had anticipated"—were impatient.

"It was all legal," Peckham says. "Whether it was ethical or not is another question. Well, I guess it depends upon what your ethics are, right?"

Wires in the warehouse

You don't have to dig into archival footage to see the Freehand in action. A few miles from where Mumford lives in the Denver suburbs, I met Scott Abram, an accountant for the U.S. Department of the Interior. Abram broke his neck in 1989, at age 17, when he dived into a shallow river on a high school field trip. He got the Freehand a decade later and still uses it for certain tasks. When we had lunch in a restaurant, he ordered a chicken sandwich. By activating the Freehand with shrugs

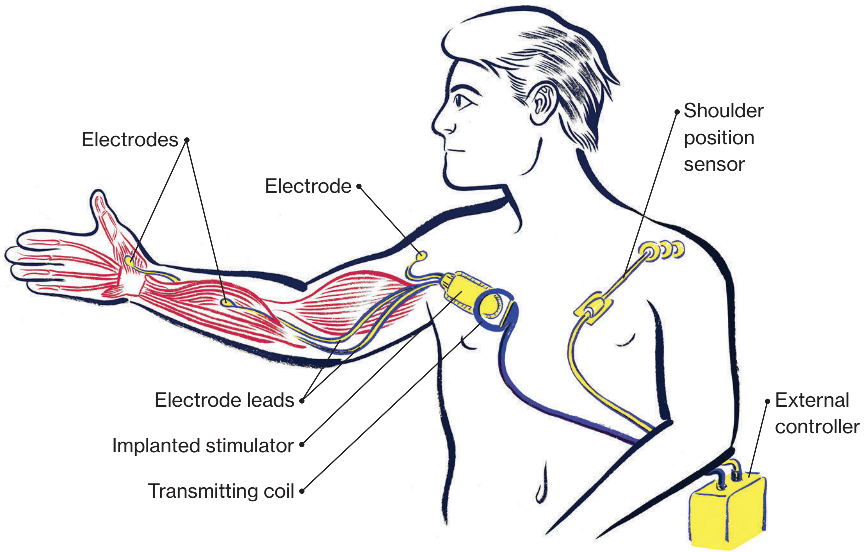

Components installed inside Freehand patients fared better than the ones on the outside. The diagram above shows how wires ran from one shoulder to a control unit and from that unit to a transmitting coil. Below, Mumford and his left hand in a 1990s promotional video for NeuroControl.

DIAGRAM BY KRIS MUKAI

of his left shoulder, he was able to manipulate his right hand in ways that helped him bring the sandwich to his mouth and down to the plate. All the while, a pager-like control unit on the left side of his wheelchair was still doing what it has done for 15 years: telling the stimulator in his chest which wires in his right arm needed jolts of electricity.

Abram knows full well what Mumford went through when the wires on the outside of his body needed to be replaced. It happens to him, too. There's one key difference, though: several years ago, Abram managed to track down Kevin Kilgore, one of the researchers who developed the technology with Peckham in Cleveland. And Kilgore has been sending him wires over the years.

The situation mystifies and upsets Kilgore as much as anyone. When NeuroControl was in business, it supplied the Freehand to surgeons who installed it and served as the patients' point of contact. From the perspective of patients like Mumford, the researchers who had originally invented the technology were not in the picture at all. When NeuroControl folded, nearly everything about it fell into a black hole. Not only did it fail to arrange technical support for its customers, but its website and phone number went out of service, leaving both the surgeons and the patients in the dark about what they might do next. Kilgore and Peckham say the company even refused to give them a list of patients who had gotten the implants. To this day the engineers say they don't know exactly how many there were.

For Damion Cummins of Monroe, Louisiana, the company's demise had a surreal aftermath. He had gotten the Freehand after being paralyzed in a high school football game. But it didn't consistently work as well as he hoped, and he stopped using it after less than two years. Stopping was easy enough—he no longer asked someone to tape the awkward external wires to the device in his chest. But as the years went on, he wondered about that dormant electrical equipment, some of which you can feel right under his skin. "Is it going to disintegrate or break off?" he asked himself. "Should I worry about that?" He thought about going to see the surgeon in Shreveport who had implanted the Freehand, but the doctor had moved to California. Cummins says he spent a few years feeling uneasy about the electronics in his body before he finally tracked down the surgeon and called him. "Should I have it taken out?" Cummins asked. "No, as long as nothing's bothering you," the doctor said.

It's painful for Kilgore to hear about the isolation that Cummins felt. About five years ago Kilgore got a $75,000 grant from Paralyzed Veterans of America, a nonprofit group,

to follow patients with electrical-stimulation implants over an extended period. He spent much of the money buying up one of the few chunks of NeuroControl that hadn't completely vanished: its inventory of wires, stimulator coils, controllers, batteries, and other Freehand parts, which another Ohio company had bought and was keeping in a warehouse. With that stockpile, Kilgore reached out to the Freehand patients he and his colleagues did know of—a few dozen people in Ohio—and set up an online users' group in hopes of finding more.

In 2009, Kilgore and other researchers tracked down 65 Freehand recipients and determined that more than half were still using the device. Today he estimates that he has enough parts to keep such patients going for a few more years. But eventually, he says, "the ultimate fix" is for the patients to get something better. Nearly 30 years after the birth of the Freehand, the Case Western team has improved the technology significantly. Among other things, they have made the control unit small enough to be implanted in the body, eliminating the need for external wires that can snag and break. The device can also do more than restore grasping ability. It can be networked, as they put it, to send electrical stimuli to many more muscles—providing upper-body support, for example, or bowel and bladder control. The researchers have gotten some paralyzed people to stand and take halting steps with the help of a walker.

The essential economic dilemma remains, though: without a company to market this technology widely, the pool of potential recipients is limited to people who live in or can afford to travel to Cleveland. And if it's not a commercial product, insurance companies won't cover the cost of the device. That means the researchers have to rely on grant money to get these technologies into patients. "I can do five implants a year on grants," Kilgore says. "But I get 100 phone calls a year."

Even hundreds of patients a year might not make for a big enough market to entice private investors. But Kilgore and Peckham think they may have figured out a solution.

Deepening the pool

They are convinced that avoiding a repeat of the NeuroControl fiasco with many future implantable technologies will require a nonprofit/for-profit partnership. They've formed the nonprofit: the Institute for Functional Restoration at Case Western. Its mission is to usher technologies through regulatory approval; after that it could market the devices itself or license them to for-profit companies. Ideally, if such a company failed, the nonprofit—funded mainly by a private foundation—could keep supporting patients.

The first technology the institute will handle will be the networked device that is the descendant of the original Freehand. The organization has grants to begin a clinical trial and even to develop a manufacturing facility for the devices. It also has a waiting list of potential patients. But it has yet to sign up any companies as for-profit partners—companies that, as Peckham puts it, are "not trying to meet some venture expectations of how fast you return their investment."

In theory, there could be many potential partners. As it happens, the neurostimulation business is enjoying a renaissance, especially in Cleveland, given the abundance of technologies to license from Case Western, the Cleveland Clinic, and other centers there. Several of the companies are staffed with alumni from NeuroControl, including Thrope, who now heads NDI, a firm that invests in neurotechnologies. Thrope says partnering with a nonprofit would be attractive to companies that don't want to bear the risks inherent in taking a new technology through years of testing and regulatory approval. If the nonprofit can handle that part and then turn things over to a for-profit company, Kilgore and Peckham's model "has some worthiness to it," he says.

But even with that risk removed, Thrope is quick to add, not a lot of companies are interested in selling products that only a small group of people can use. Instead, he says, he and other investors are eager to find opportunities to address what doctors call multiple "indications," meaning they can treat more than one condition. He mentions Second Sight, a publicly traded maker of a $140,000 retinal implant that can restore sight to people with a hereditary form of blindness. The potential market is quite big—perhaps 1.5 million people worldwide and 100,000 in the United States—but even so, Second Sight is already testing ways to deepen the pool of patients by treating other forms of blindness. Thrope says his firm, which he founded in 2002, rarely jumps in to invest in a neurotechnology until it has been developed beyond its initial stage and can treat a second or third indication. It's "reversing the formula we used in NeuroControl," Thrope says. "We've tried to avoid breakthrough technologies if possible."

Avoiding breakthroughs: that seems to go against our tendency to imagine that technology will fix so many broken things, our bodies included. But consider the perspective of Damion Cummins. He says he endured multiple surgeries to get the Freehand because anything that could improve his daily life was worth a shot. He accepted the idea that it might not work. But when I asked him if he would have gotten the implant if he had realized there was a chance NeuroControl could fold, he replied: "If I had known that, then I definitely would not have."

Brian Bergstein is deputy editor of MIT Technology Review.


How can you make batteries love dancing?
The demand for energy is rising fast. To help the world keep up, we are celebrating our 150th anniversary by co-creating smarter ways to harvest and generate energy, like using the footsteps on dance floors to charge clubbers' phones.
Join the conversation at creator-space.basf.com
150 years
BASF
We create chemistry

The time is now.

This October, extraordinary people will be inspired to work together on the world's hardest problems, improving the way we all Learn, Cure, Fuel, and Make.

SOLVE

Providing quality education to anyone who wants to learn, anyplace.

Leveraging innovation to make health care universally available.

Generating clean energy and water, and enough food to power the planet.

Creating infrastructure and opportunity for 10 billion people.

Let's Begin | **October 5 - 8, 2015** | The MIT Campus

Massachusetts Institute of Technology

HUB week
October 3-10, 2015
From curiosity, creation.

solve.mit.edu

MIT Technology Review
BUSINESS REPORT

Path of Persuasion

How technologies from smartphones to social media are used to influence our tastes, behavior, and even habits.

CONTENTS

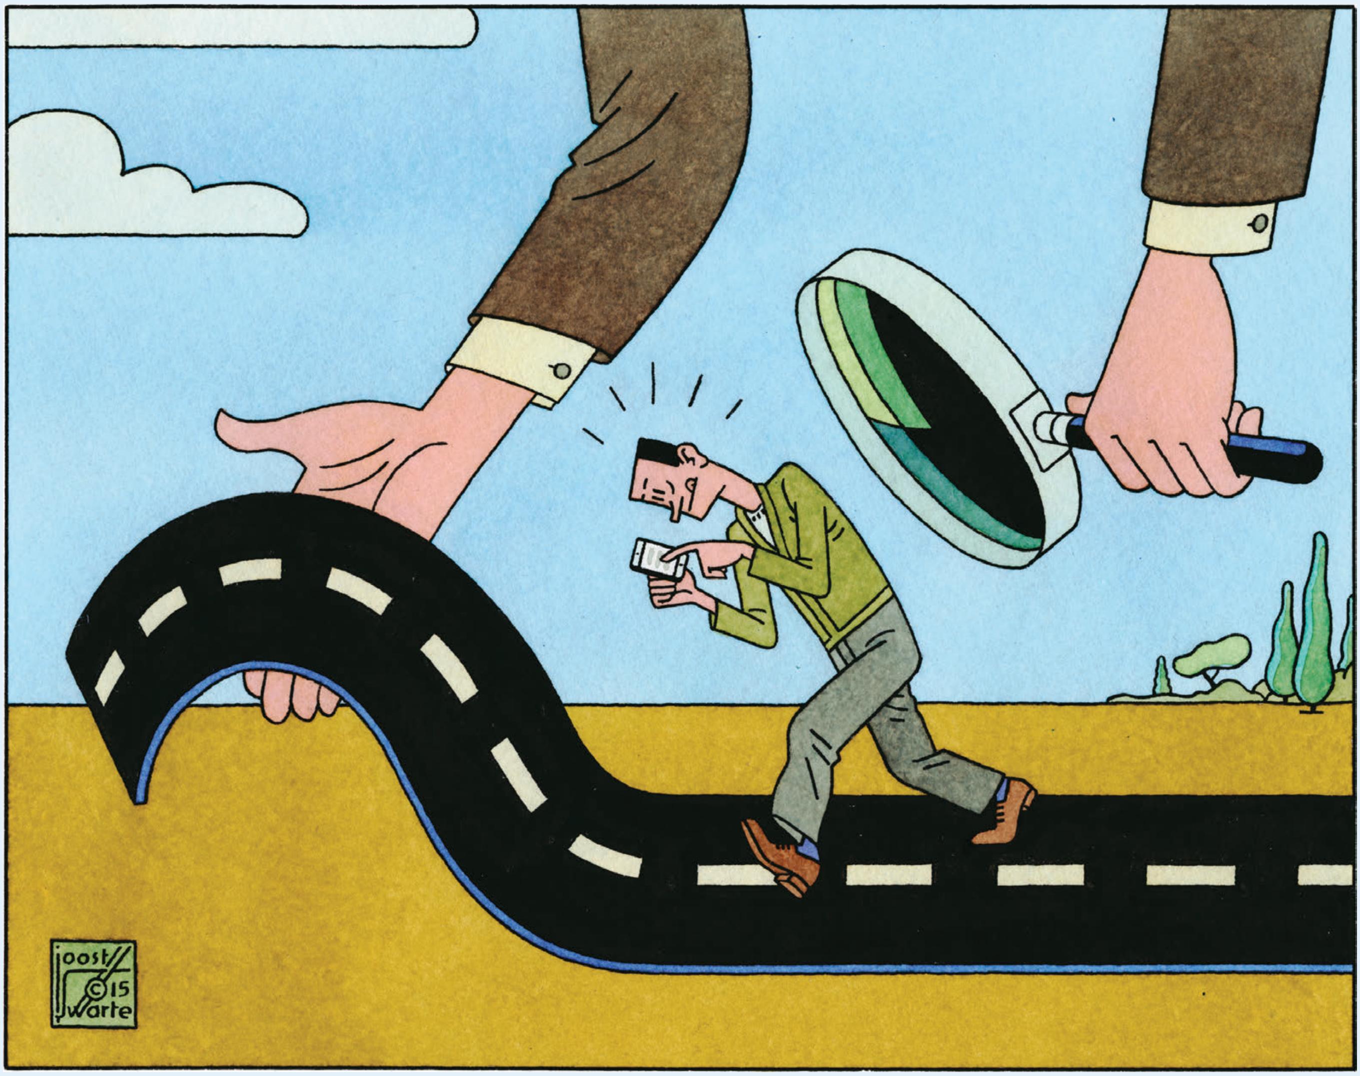

JOOST SWARTE

The Big Question

Technology and Persuasion

Persuasive technologies surround us, and they're growing smarter. How do these technologies work? And why?

● GSN Games, which designs mobile games like poker and bingo, collects billions of signals every day from the phones and tablets its players are using—revealing everything from the time of day they play to the type of game they prefer to how they deal with failure. If two people were to download a game onto the same type of phone simultaneously, in as little as five minutes their games would begin to diverge—each one automatically tailored to its user's style of play.

Yet GSN does not simply track customers' preferences and customize its services accordingly, as many digital businesses do. In an effort to induce players to play longer and try more games, it uses the data it pulls from phones to watch for signs that they are tiring. Largely by measuring how frequently, how fervently, and how quickly you press on the screen, the company can predict with a high degree of accuracy just when you are likely to lose interest—giving it the chance to suggest other games long before that happens.

The games are free, but GSN shows ads and sells virtual items that are useful to players, so the longer the company can persuade someone to play, the more money it can make. Its quickly growing revenue and earnings are a testament to how well this strategy works, says Portman Wills, GSN's chief information officer. Along with factors such as smart engineering and creative design, using data to shape persuasive tactics is a key to the company's success.

The idea that computers, mobile phones, websites, and other technologies could be designed to influence people's behavior and even attitudes dates back to the early 1990s, when Stanford researcher B. J. Fogg coined the term "persuasive computing" (later broadened to "persuasive technology"). But today many companies have taken that one step further: using technologies that measure customer behavior to design products that are not just persuasive but specifically aimed at forging new habits.

If habit formation as a business model was once largely limited to casinos and cigarette manufacturers, today technology has opened up the option to a broad range of companies. Insights from psychology and behavioral economics about how and why people make certain choices, combined with digital technologies, social media, and smartphones, have enabled designers of websites, apps, and a wide variety of other products to create sophisticated persuasive technologies.

How these technologies work and why are the big questions this Business Report will answer.

Insights from psychology and behavioral economics about how and why people make certain choices, combined with digital technologies, social media, and smartphones, have enabled designers to create sophisticated persuasive technologies.

With new digital tools, companies that might once have been simply hardware makers (such as Jawbone) or service providers (Expedia) are now taking on the role of influencer, attempting to shape the habits of their users by exploiting the psychological underpinnings of how people make choices.

While Expedia is trying to design its website so as to trigger someone to visit daily, Jawbone has built features into its fitness bands and other products that executive Kelvin Kwong grandly describes as "using our best understanding of how the brain works to get you to act." And Kwong says it's working. Sending carefully designed messages to people wearing Jawbone fitness trackers has helped them get an additional 23 minutes of sleep per night, on average, and move 27 percent more, the company says.

Habit Design, which bills itself as "the leading habit training program," employs game designers and people with PhDs in behavioral science. It says it has created a platform that keeps 80 percent of participants in corporate wellness programs involved over three months. Traditional programs like seminars or counseling, by contrast, generally lose 80 percent of participants in the first 10 days, according to Michael Kim, a former Microsoft executive who is now Habit Design's CEO.

New data-centered models of persuasion are having an impact not only on new startups but on traditional influencers, from political consultants to advertising agencies. In politics, data consulting firms that emulate the kind of voter modeling, mobilization, and persuasion the Obama campaigns pioneered are multiplying. One model for today's new type of ad firm is Rocket Fuel, based in Redwood City, California. Staffed by people with PhDs in game theory and predictive modeling, the firm uses artificial intelligence to predict the best ad to show a given customer looking at a particular Web page, taking into account data gathered from websites; the browsing, advertising, and purchase history associated with a given shopper's IP address; and insights into what style of ad works best on a certain website (blue hues are best on Answers.com, for example). Founded in 2008, the company claims its targeted ads generate revenue for clients amounting to two to eight times what is spent on the ads. Last year Rocket Fuel had revenue of more than $400 million.

Marketers argue that there's potential for all this to benefit consumers, who want better service and more suitable offers. "They expect companies have data on them. They just want it to do something useful for them," says Philip Wickline, CEO of Zaius, a Boston-based startup building a platform that will allow a company to track customers' behavior, with their permission, as they interact with it

in stores, online, and in any other context. Armed with this information, companies could better understand the value of each customer and more effectively measure the return on ads or discounts directed at that person.

Given the depth of information about us that tracking technologies generate, and companies' increasingly sophisticated attempts to affect our behavior, what are the appropriate limits of a kind of persuasion that can be so well designed as to be nearly invisible? There are already legal limits on how companies can advertise products. But the government's own use of behavioral persuasion has led to calls for updated regulations.

Rather than trying to regulate hard-to-spot attempts to get people to form new habits, a more practical solution might be for product designers to agree to adhere to principles like transparency and disclosure. Requiring a user to sign up to be persuaded—as you very well might in search of better sleep or fitness—could be best. —*Nanette Byrnes*

Persuasive Texting in Mozambique

Inspired by trials that used text messages to change behaviors, a U.K. charity tested whether it could persuade HIV-positive people to attend their appointments.

PROBLEM About 1.6 million people in the southeast African country of Mozambique live with HIV. Antiretroviral therapy can prevent its spread, but only 74 percent of patients who start HIV treatment are still taking the medicines a year later.

SOLUTION In November 2011, the U.K.-based children's organization ARK began a two-year test of sending text messages to HIV-positive people in urban and rural areas of Mozambique's Maputo Province to remind them about treatments and appointments. About 15.5 million mobile phones are already connected in the country of 25.8 million people, and use of the technology is growing.

STUDY ARK's study followed 830 men and women undergoing antiretroviral therapy for an average of 16 months, and 522 HIV-positive expectant mothers until eight weeks after giving birth. Patients were randomly assigned to receive the text reminders or to continue regular treatment with no reminders.

METHOD Using a database of electronic patient records, a computer program sent out automated text reminders at times including a week before an appointment, two days before the appointment, and two days after any missed appointments. HIV-positive pregnant women also received educational messages about testing and treatments.

RESULTS The SMS messages helped some patients improve their treatment regimen, but not all. The messages significantly helped urban and recently diagnosed HIV patients stick with treatment. Such patients who didn't receive text messages were nearly twice as likely to fall off the program, failing to consistently pick up drugs or attend appointments. Results at the rural centers were disappointing, possibly because of transportation issues, limited cellular coverage, migration to other provinces, and a limited sample size.

The texts did help persuade mothers to test their newborns for HIV, but overall the program didn't make them more likely to complete pre-birth treatment or to give birth in a health center.

NEXT Now that texting has been shown to help some patients, several groups are launching a much bigger messaging program in another Mozambican province. This program could reach 58,000 patients by the time it concludes in 2016. Researchers will examine how cost-effective the trial is, how to improve technology in rural areas, and how best to support patients who are most at risk of not completing treatment.

Q&A

New Technologies Persuade in Old Ways

Robert Cialdini, an expert in the science of persuasion, talks about its most modern methods.

● For anyone interested in the science of persuasion, psychology professor Robert Cialdini has been the expert of choice since the 1970s. During his long career at Arizona State University, he has studied everything from the ways blood banks attract donations to the reasons why some people pick up litter and others don't. Cialdini's best-known book, *Influence*, has sold 2.5 million copies worldwide. It remains one of Amazon's top 500 titles, 31 years after its debut.

Cialdini argues that practically every form of persuasion can be traced back to one of six timeless principles: reciprocity, likeability, authority, scarcity, consistency, and social proof. It's human nature to reach for those levers, and to be influenced by them, he contends. *MIT Technology Review* contributing writer George Anders visited Cialdini to discuss what's changed—and what hasn't—as today's influence peddlers use rapidly evolving technologies to ply their trade.

One of your core principles is that if you want to persuade people, it helps to be likeable. Is that still relevant if we're doing most of our business online?

Yes. People often claim that e-mail and the Internet are bloodless. But there are lots of ways to put the humanity back in. There's a piece of research showing that if you include a cartoon with your opening e-mail, you get a better outcome in negotiations. You're not an opponent anymore. You're a person who has a sense of humor, and who wants to make this a

positive experience. In fact, if you mimic another person's writing style, that gets better negotiating outcomes too. If the other person uses emoticons, you do too. Or exclamation marks.

You've also said for a long time that we heed authorities' advice. Dentists telling us what brand of toothpaste to buy, for example. Today do you need an advanced degree to be an authority?
Bloggers and product reviewers are the acknowledged experts now. I've read that 98 percent of online shoppers read product reviews before deciding what to buy. That's a mind-blowing statistic. We can't get 98 percent of people to agree that the earth is round. But now you can find authorities in Turkey or Tokyo or Portland, Oregon, that we didn't have access to before.

We've always valued social proof. Now the Web has given us access that we never had before to large communities of individuals just like us. That's revolutionized marketing. Look at Kiddicare, which is an online superstore for baby products. They categorize ratings in all sorts of ways: first-time parents, parents of twins, and so on. You're going to go to these intensely comparable voices.

In the old days, store executives tried to impose their priorities on you. Kmart announced blue-light specials in hopes that shoppers would stampede toward that part of the store. That's not the best approach anymore, is it?
It's much more trustworthy when you tell customers what their peers are saying. I trust customers. I don't trust the store manager.

Isn't there a risk that we could spend so much time on an island of like-minded people that we never get to see the rest of the world?
I worry about this. I think you're finding it in political attitudes that are becoming more polarized. We tend to go only to those places where we hear opinions that we want to hear. We end up reinforcing our preconceptions.

You're a big believer in the importance of scarcity—or at least the appearance of scarcity—as a way of getting customers excited. On the Internet, supply and choices seem infinite.
People can create scarcity with ... an offer available for only a short time. Groupon does it, and quite expertly, too.

Or restaurants that offer one table per night via online reservation sites?
Exactly. It's price be damned. It's hardly a bargain that you're trying to get. It's a distinctive, rare experience.

I'm thinking of Taylor Swift, who responds to lots of her fans via social media. She's the one singer who still can sell a million CDs. In this case, it's not scarcity at work so much as another of your principles: reciprocity. Fans feel duty-bound to buy her music in a pricey format rather than settling for cheaper downloads or free bootlegs. Is that because she's reached out to them?
That's right. You just don't feel that you can take advantage of someone who has been a benefactor.

You've written about "the unthinking yes"—in which these persuasive techniques sweep us along so smoothly that we agree to do things without even reflecting on how we got there.
What's good about automatic responses is that they give us a lot of efficiency. We don't have to think very much about decisions. We will still do well if we follow social proof, or an authority. But the downside is that we're not thinking. We're reacting. And it does make us vulnerable to people who counterfeit the evidence.

That said, it's inevitable. The pace of information is accelerating. We don't really have the luxury to step back and think hard about the great majority of decisions that we make every day. And so we need these shortcuts.

Pretrial Technology

Lawyers are testing arguments and evidence online.

ACADEMIC ORIGINS Borrowing a tool from social scientists, academic researchers studying jury dynamics now run tests on groups pulled from the online marketplace Amazon Mechanical Turk. One study by Jessica Salerno, an assistant professor at Arizona State University, found that Mechanical Turk mock jurors were more likely to vote guilty when the graphic crime photos shown them were in color.

IMITATION The trial consulting firm DecisionQuest has built its own Web tool with 3.5 million mock jurors. A legal team can use it, for example, to test the impact of different arguments.

TIME In 24 to 48 hours, online jury research can provide the same number of responses that traditional jury research provides in one to two months. Lawyer Michael Cypers recently used online sampling to help a client decide whether to take a settlement offer with a 10-day time limit. It was helpful, he says, to know that the online jurors were skeptical of the fraud claim against his client and not convinced that negligence was involved.

COST An online mock jury exercise can cost \$5,000 to \$20,000, whereas a live mock jury costs \$50,000 to \$150,000. Mechanical Turk workers will take a 20-minute survey for as little as \$5 per person, explains lawyer Jonas Jacobson, a trial preparation expert.

IMPACT The Focal Point, a firm specializing in the visual presentation of cases, sometime uses online testing to evaluate whether evidence like simple diagrams or 3-D animations is well-recalled and useful in mock deliberations.

LIMITS Cypers says that while this testing is useful, it won't replace testing with live mock jurors, which offers the opportunity to "look in their eyes and see how they respond."

Technology

Fake Persuaders

Fake accounts can inflate follower counts, suppress political messages, and run stealthy social marketing.

● Advertising revenue is soaring at Facebook and Twitter as consensus grows that people can profitably be influenced by promotional messages woven in between updates from their friends.

But not every commercial enterprise exploiting the persuasive power of social media has set up a corporate account or pays for ads. Fake accounts operated by low-paid humans or automated software have become good business, too. They are used to inflate follower counts, to push spam or malware, and even to skew political discourse. The tactic appears to be pervasive and growing in sophistication.

On Twitter as many as one in 20 active accounts are fakes. Facebook's equivalent number is a little more than one in 100 active users. Software tools that help you make new accounts in bulk can easily be found or bought online, says Christo Wilson, an assistant professor at Northeastern University who has studied the problem of fake accounts.

One of his students recently tested some of those tools and set up 40 Twitter accounts and 12 Facebook accounts in a single day before the companies blocked new registrations from that Internet connection. Simple evasive measures would probably have allowed many more accounts to be made. Investors closely scrutinize active user counts to gauge the value and potential of social networks. That encourages sites to ensure that their security systems don't block legitimate users, says Wilson, making it easier for fake accounts to flourish.

Fake accounts are given a veneer of humanity by copying profile information and photos from elsewhere around the Web. They can gain fake friends by exploiting human nature and the fact that people on a social network are often looking for new connections and content. "Choose a picture of a beautiful woman, and all of a sudden people accept your friend request," says Wilson. Celebrities often have large numbers of fake followers because aping what many real users do is an easy way to make a fake account look legitimate.

Once a fake account is established, the simplest way to make money with it is by quickly inflating the numbers of things like followers or "likes." It is easy to find sites offering 100,000 new Twitter followers for as little as $70. Instagram and Facebook "likes" and Pinterest "pins" are also easily bought. Having more followers or likes helps people and businesses look good. It can also influence the algorithms used by social networks or other companies to recommend influential accounts.

Fake accounts have been used in more sophisticated ways to fake social support for something, and to influence real users to join in. The accounts are controlled either by software or by paying Internet users in developing countries a few cents per action.

In 2010, a conservative group in Iowa used automated accounts to send messages supporting Republican candidate Scott Brown's attempt to win a Massachusetts seat in the U.S. Senate. Thanks to retweets by some real users, the messages reached an audience of 60,000. In Mexico's 2012 general election, the Institutional Revolutionary Party used more than 10,000 automated accounts to swamp online discussion. Both parties won their races, although it's not clear what impact these social-media manipulations had.

Recently, automated accounts have been seen staging more commercial campaigns. A 2014 study of 12 million users of China's influential Weibo social network, which is similar to Twitter, found 4.7 million accounts involved in campaigns that try to manufacture word-of-mouth support for particular products. Most were automated accounts that amplified certain messages, mentioning products or services, from people with large followings (messages likely paid for by the brand behind them). Also last year, automated tweets were part of a scam that inflated the value of penny-stock tech company Cynk to $5 billion in just a few days.

Filippo Menczer, a professor at Indiana University, says more sophisticated "social bots" that engage with other users are probably active on Twitter and other networks but escaping detection. Research experiments with such bots have shown that they can successfully gain social capital and even shape the social connections humans make with one another, says Menczer.

As social networks become more tightly coupled to personal spending and wider economic activity, the incentives to use them grow stronger, Menczer says.

As social networks become more tightly coupled to personal spending and economic activity, incentives grow.

In 2014, the security company Bitdefender picked up a social bot using names including "Aaliyah" that was stalking men on the casual-dating app Tinder. Aaliyah would start a simple, scripted conversation, then ask the victim to play a particular social game, offering her phone number in exchange. The scam didn't have a clear business model, but Bogdan Botezatu, a senior threat analyst at Bitdefender, believes it was "a test run for something much bigger."

The Pentagon's research agency DARPA, which has its own concerns about what it calls "deception or misinformation campaigns" in social media, sponsored a contest in which teams of researchers compete to detect social bots at work in a Twitter-style social feed. Menczer, who took part, hopes the contest will lead to tools that are better at policing real social networks. "It's kind of scary that we don't know how to detect these kind of bots and campaigns if they are out there," he says.

—Tom Simonite

Profile

Compulsive Behavior Sells

Nir Eyal is showing software designers how to hook users in four easy steps. Welcome to the new era of habit-forming technology.

● A middle-aged woman sits before a computer screen on the 11th floor of Expedia's glass-clad headquarters in Seattle. Two electrodes are taped to her brow just above her left eye, two more on her left cheek. A one-way mirror reflects her face as she responds to requests issuing from a speaker mounted in the ceiling.

Behind the glass, a researcher directs the test subject as a half-dozen designers, engineers, and executives look on in rapt silence. "Okay, Shannon," the researcher says. "Go to Expedia and start shopping for your trip to Hawaii." The audience gazes intently at a large video display. A running graph of the electrodes' output trails across the screen. The electrodes on the brow measure contraction of the muscles that activate frowning—a sign, according to the theory of facial electromyography, of concentration, tension, or irritation. Those on the cheek track the play of muscles involved in smiling, evidence of the warm glow of delight that occurs when the brain's reward circuitry is activated.

Though she has not been told, Shannon has been brought in to test a new Expedia feature, known as Scratchpad, that the online travel broker hopes will bring travelers back to the site daily between the time they start planning a trip and the day they make a purchase. Scratchpad automatically records the hotels and flights a customer has viewed, allowing users to pick up on previous searches without having to re-create them.

The first part of Shannon's test is simply finding the Scratchpad button on the Expedia.com page. She hasn't found it yet, but when she looks at photos of the Westin Maui, the smile sensors spike: a jolt of joy—and potential paydirt for Expedia. The company aims to make the experience of shopping so pleasurable that using the site becomes a habit.

Forging new habits has become an obsession among technology companies. In an age when commercial competition is only a click away, the new mandate is to make products and services that generate compulsive behavior: in essence, to get users hooked on a squirt of dopamine to the brain's reward center to ensure that they'll come back.

The rise of mobile computing has intensified that imperative. The small screen crowds out alternatives, focusing a person's attention on a limited number of go-to apps. The ones that get used are the ones people click on impulse while they're drinking their morning coffee, waiting for the bus, or standing in the checkout line.

For a long time, the methodology for designing habit-forming products was haphazard: build it, put it before the public, and watch it go viral or fade into oblivion. In recent years, though, product teams have become more deliberate. Principles derived from behavioral science play an increasing role in software design, creating a demand for experts who can guide developers in the art—and science—of behavior engineering.

"The confluence of increased access and greater sharing of personal information, and at higher transmission speeds, has created the perfect storm of addictive technology." —Nir Eyal

Among the most influential is Nir Eyal, an entrepreneur turned user-experience guru who has become Silicon Valley's most visible advocate of habit-forming technology. His blog, Nir and Far, has attracted more than 25,000 subscribers hungry for insights into shaping user behavior, and his writing has appeared in both the mass-market pages of *Psychology Today* and the insider club of TechCrunch.

He has worked for some of the biggest names in technology (most of whom don't want to talk about it) and has presented workshops from Norway to Thailand. His inaugural Habit Summit, held last March on the Stanford campus, drew hundreds of participants from startups and blue-chip firms alike. Eyal promotes a scheme he calls the hook, a simple set of steps derived from his observation of numerous online products and services and undergirded by a wide range of psychological and neurological research. The hook, he says, is the magic behind Facebook, Google, Instagram, Pinterest, Snapchat, Twitter, and just about every other icon of the consumer Internet. It leads users into a repetitive cycle that transforms tentative actions into irresistible urges.

John Kim, Expedia's chief product officer, brought Eyal in last year to help the company develop compulsive experiences, and now Shannon, the volunteer in front of the computer screen, is testing the fruit of their labor.

After several minutes, she still hasn't discovered the Scratchpad feature. Finally the researcher guides her to its unobtrusive button in the menu bar at the top of the Expedia homepage. Low-level tension registers on the graph at the front of the room—concentration? frustration?—but then she recognizes a picture of the Westin Maui, which Scratchpad captured automatically. A shimmer of delight ripples across the graph. "I like the fact that it saves me time," she comments. "I'd go back and use this again for sure!"

Eyal's workshops offer a four-hour immersion in the mechanics of the hook. On a warm spring day several weeks before the Expedia research exercise, he's getting ready to run a session at the office of Zurb, an airy design studio not far from Apple's headquarters in Cupertino, California. Clad in a plaid shirt that hangs untucked over jeans, he's coat-hanger thin with eyes that glitter beneath a clean-shaven crown.

He opens with a disclaimer. "I'm not an advocate for creating addiction," he says. "Addiction has a specific definition: it always hurts the user. I talk about the pathways for addiction because the same things that occur in the brain help us do something that can be good."

Thus he initiates 67 attendees from companies including Hewlett-Packard, the New York Times, and Samsung into the mysteries of the hook.

It starts with a trigger, a prod that propels users into a four-step loop. Think of the e-mail notification you get when a friend tags you in a photo on Facebook. The trigger prompts you to take an action—say, to log in to Facebook. That leads to a reward: viewing the photo and reading the comments left by others. In the fourth step, you inject a personal stake by making an investment: say, leaving your own comment in the thread. This pattern, Eyal says, kicks off a cycle that lodges behaviors in the basal ganglia, the part of the brain where automatic behaviors are stored and where, according to neuroscientists, they last a lifetime.

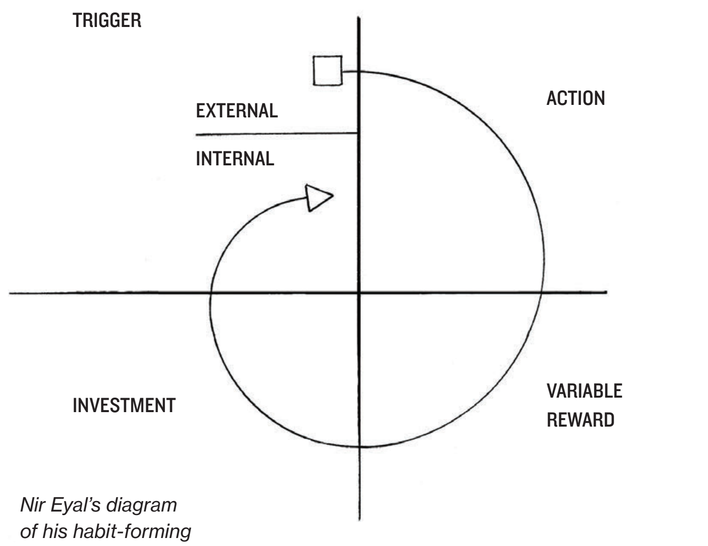

Nir Eyal's diagram of his habit-forming "hook."

The psychology behind the hook dates back at least to the 1930s, when the American psychologist B. F. Skinner showed that he could induce desired behaviors in animals. Skinner is famous for training pigeons to do seemingly intelligent things, like reading signs and following instructions by manipulating the equivalent of Eyal's trigger-action-reward sequence.

Other researchers have refined Skinner's theories in the decades since. On the screen behind him, Eyal flashes a slide drawn from the seminal work of Stanford behavior theorist B.J. Fogg. It's an x-y plane with axes labeled "motivation" and "ability," a curve tracing a diagonal smile from upper left to lower right. According to Fogg, a behavior happens when a trigger coincides with both motivation and ability—but only in the right proportion. If a trigger consistently fails to initiate the desired action, the theory goes, habit designers should aim to enhance the user's ability. Motivation is hard to influence, because you can't make people do what they don't want to do. Ability is more malleable: simply make the behavior easier to execute.

Still, the reward must promise enough pleasure to drive people to take the intended action. In training animals to execute complex behaviors, Skinner discovered that varying the payoff—from highly desirable to nothing at all—both increases a behavior's frequency and helps keep it from fading once the rewards stop.

A classic example is slot-machine gambling. The player never knows whether the next pull might bring a $5 win or a $50,000 jackpot. The unpredictability of the reward—and the randomness of its arrival—is a powerful motivator to pull the lever again and again.

Eyal draws a parallel between Skinner's variable rewards and the endless variety to be found on, say, Pinterest: the user can scroll endlessly, scanning for distinctive items amid a sea of banality. "This," he says, swiping his finger downward as though scrolling a touch screen, "becomes this"—he moves his arm up and down as though he were cranking a slot-machine lever.

The hook's final stage, investment, closes the loop by "loading the next trigger," Eyal says, an idea inspired in part by work on game psychology by Jesse Schell, a Disney Imagineer turned Carnegie Mellon professor. Take Twitter. When you make an investment by posting a tweet, a follower's reply to your contribution triggers an e-mail notification to your inbox, inciting you to take yet another spin through the cycle.

The workshop hums with activity as the students form small groups to work on their own projects. One of the participants is an Expedia executive named Pooja Vithlani, who is part of the team developing Scratchpad. (Her title is senior product manager of compulsion.) John Kim sent her to learn more about how to apply the hook.

She has a clear sense of the Scratchpad user's action (shopping for airline tickets), reward (a handy list of possible travel arrangements), and investment (curating the list by eliminating options that prove obsolete or impractical). However, the trigger proves elusive. Vithlani muses on the anxiety occasioned by oncoming holidays and the attendant travel plans. Can she nudge potential customers to check Scratchpad at the very moment they're feeling pressure to lock in airfares? Relief from the stress of holiday travel might be the bait on Scratchpad's hook.

Eyal readily admits that many of the ideas he promotes aren't his own. "I don't do original research and don't intend to," he says over his habitual lunch, a build-it-yourself burger at an eatery near his home in Palo Alto, California. "There's more research than we know what to do with."

To his clients, he offers a simple, practical scheme that keeps the arduous process of product design focused squarely on

The Roots of the Hook

Nir Eyal's methodology draws from earlier research.

1930s
Psychologist B. F. Skinner induces desired behaviors in animals.

1957
Skinner and behavioral psychologist Charles Ferster publish *Schedules of Reinforcement*, which describes experiments in which pigeons are shown an "antecedent stimulus" such as a colored light and then given rewards of food when they exhibit a desired behavior, such as pecking a spot on the wall, at the same time. In this manner the birds are trained to peck that spot whenever they see the light. The book contains a key insight: that varying the payoff for a behavior both increases a behavior's frequency and keeps it from fading once the rewards stop coming. Think slot machines.

1963
Psychologist Albert Bandura, with Richard Walters, publishes his social learning theory, which holds that our actions are shaped not only by anticipated consequences but also by what we see other people doing. Eyal will view this as a force in social networks.

1985
Psychologists Edward Deci and Richard Ryan introduce self-determination theory, which says that people are driven by three innate psychological needs: autonomy, competence, and relatedness. Eyal will see this theory, especially the need for autonomy, at work in video games and online reviews.

1990s
In monkey experiments, neuroscientist Wolfram Schultz shows that cues preceding behavior that is then rewarded set the brain on alert for further cues. These cues are what Eyal later calls external triggers. Eyal will argue that effective external triggers form an association in the mind with internal triggers like loneliness or fear.

2002
Psychologist B. J. Fogg coins the term "persuasive technology." In 2009 Fogg will publish a model that describes behavior as the result of a trigger coinciding with motivation and ability, in the right proportions.

2010
Game designer and theorist Jesse Schell gives a talk, "Beyond Facebook," at a conference in Las Vegas. The presentation helps spur Eyal's recognition of the role of personal investment in cementing habits.

2011
Behavioral economist Dan Ariely and coauthors Michael I. Norton and Daniel Mochon, both professors of marketing, publish their "Ikea effect" research, experiments showing that people value their own creations more highly than others value them. We place a value on the sweat equity we put into creating a YouTube video or a custom playlist on SoundCloud.

2012
New York Times reporter Charles Duhigg publishes the bestseller *The Power of Habit: Why We Do What We Do in Life and Business*. The book tells readers how to rewire existing behavior using a Skinnerian three-step process of cue, routine, and reward that he calls "the habit loop."

2014
In his book *Hooked: How to Build Habit-Forming Products*, Eyal adds a fourth step—investment, which "loads the next trigger" to form a feedback loop—to Duhigg's basic scheme. His premise is that people grow attached to a behavior when they invest their own effort in it.

a user's impulses, desires, and motivations. And there's a chance he'll come cheap. He charges some clients a day rate determined by rolling a pair of dice and multiplying the result by 100. "It's a variable reward," he says with a sly grin.

Eyal developed his interest in habits early. Born in Israel, he moved with his family to Florida at age three. Eating provided the surest relief from the alienation of being a foreigner with a funny name. "My parents took me to a fat camp when I was 12," he says. He racked up Cs and Ds despite being pegged as a gifted student.

In adolescence he began to shed the pounds and focus on his studies after reading *The T-Factor Diet*, which emphasized methodology over willpower. The experience showed him how much unconscious impulses influenced his own life, and the power to be gained by working with rather than against them.

Eyal graduated from Emory University with a degree in journalism and eventually landed at Stanford's Graduate School of Business. After Mark Zuckerberg spoke to the class, "overnight everyone was making an app," he says.

He realized that all those apps would need a way to generate revenue. In 2008, he persuaded the prominent venture capitalists Kleiner Perkins and Mike Maples to fund a company, AdNectar, that would broker the coming wave of socially driven brand messaging. He sold the business three years later to social-commerce website Lokerz for a sum that neither party will disclose.

The sale set Eyal adrift. The time, effort, and anxiety of running a startup had weakened his closest relationships and added 15 pounds to his lanky frame. Slowly it dawned on him that as mobile devices got smaller, screen real estate shrank and habits became more important. He realized that habits could be at the heart of his next business—and his own rejuvenation.

Throwing himself into studying consumer psychology, he devoured research on how products influence behavior. The successful apps he had encountered through AdNectar, he noticed, had in common a cyclical feedback loop of user behavior.

Eyal's fascination with the mechanics of habit-forming technology coincided with a dramatic rise in the Internet's potential to influence behavior, enabling software developers to manipulate many of the behavioral dynamics Skinner and other researchers had identified. Smartphones became a ubiquitous channel for delivering triggers, while apps reduced complex actions to the simple press of a button. The social Web delivered a panoply of interpersonal rewards. Game designers began to speak of forming a "compulsion loop." Entrepreneurs went from tracking monthly active users to a new measure called "compulsion rate," the percentage of users who returned from day to day.

Riding a cresting wave, Eyal developed his ideas on his blog. Eventually he decided to collect his writing into a coherent form. "I set out to write a 30-page document," he says, "and I ended up with a 256-page book."

As the manuscript neared completion, Eyal employed the hook to galvanize his audience. He sent an e-mail (the trigger) to his blog subscribers inviting them to read the book in progress and critique his work. Nine hundred people responded. They received the reward of reading the advance copy, then made the investment of adding their comments on Google Docs. Eyal promised to include their names in the final product, loading the trigger to buy a printed copy and post a review on Amazon.

The scheme worked. Eyal released the self-published edition of *Hooked* in early January 2014. Within a week and a half, he had racked up 125 reviews, and the book was lodged at the top of Amazon's ranking of product design books.

It's the afternoon following Shannon's session, and the Expedia team has gathered in a conference room to decide on the next steps. The test subject understood intuitively what actions to take with Scratchpad. The facial sensors indicated that she experienced an appropriate psychic reward. She made her investment in a provisional form; Scratchpad added hotels she looked at to the pinboard automatically. (The ability to curate the list is on the drawing board for a future update.)

But the trigger remains a glaring issue. Shannon had to be told to click the Scratchpad button. Without a cue, she didn't even know it was there.

Data has shown that using Scratchpad doubles Expedia.com's compulsion rate—in other words, visitors who register for Scratchpad and use it are far more likely to return to the site within 24 hours than those who don't. But only a small percentage of visitors actually use it. More effective triggering could go some distance toward fixing that.

Vithlani has an idea that she thinks might do the trick. She calls it continuous shopping. Travel shoppers are often paralyzed because prices shift constantly, she notes. Expedia loses sales because people forget what they've found in past searches. "Continuous shopping will give them perfect memory and fresh prices," she says.

The Scratchpad window would open automatically when a visitor arrived at Expedia.com and offer to track price changes in return for registering. Then prospective customers would receive a daily e-mail telling them whether prices had risen or fallen; they would need to click through to Scratchpad to see the details. "This could develop a compulsion loop because we're getting people used to the fact that we remember what they looked at," she explains.

It would take only a few weeks to implement and test. (A year later, variations of Vithlani's idea are live on the website, and Expedia reports that using Scratchpad now triples the compulsion rate and doubles the repeat-visit rate.)

Eyal himself is not immune to the siren call of behavior engineers. In an article he wrote for *Forbes* entitled "Strange Sex Habits of Silicon Valley," he candidly laments the impact of mobile devices on intimacy between himself and his wife. Not long ago, upon slipping into bed at night, he often found himself reaching for his tablet rather than his spouse. Drawing on B.J. Fogg's behavior model, he broke the spell by pushing his late-night browsing down the ability scale. He installed a timer that turned off his Wi-Fi router at bedtime, forcing him to switch it back on before he could satisfy his craving for a late-night online fix.

The gambit improved his sex life, but the larger issue remained, he wrote: "The confluence of increased access and greater sharing of personal information, and at higher transmission speeds, has created the perfect storm of addictive technology."

Eyal's worry isn't idle. Slot-machine designers are renowned for inducing behaviors that resemble addiction, and in 2011 the American Society of Addiction Medicine began defining addiction in terms of behaviors that activate the brain's reward circuitry, whether or not substance abuse is involved.

"The ethics of this have still to be worked out," says Chris Nodder, author of the archly titled user-experience manual *Evil by Design*.

Hooked concludes with a chapter on ethics that directs behavior engineers to focus on applications that improve users' lives and that the engineers themselves find helpful. On the whole, though, Eyal views behavior engineering as a grand opportunity. "Wouldn't you like to want to exercise without thinking about it?" he asks. "Or save money every day by being more frugal? That's what this technology makes possible."

In any case, neuroscience suggests that eliminating the potential for addiction would require eliminating pleasure itself. "In the brain, our pleasure center and memory center are in close proximity, as though nature wanted us to reproduce and remember how," explains Howard Shaffer, a Harvard psychiatrist.

A future of smart watches and biometrics may make engineering new habits even easier. "Now the interface disappears, which provides all kinds of new triggering opportunities," says Eyal. "I think we'll see a golden era—I hope—of habit formation and interesting ways to help people live better lives."

—Ted Greenwald

ACCESS THE FULL REPORT ONLINE

Read more about the industries being changed by persuasive technologies and the technologies themselves.

New Marketing

Voter Algorithms

Health Persuasion

Neuroscience Advertising

Recommendation Engines

Upcoming Events

technologyreview.com/business

Reviews

Photographs by Jeff Minton

Left, Ethan Brown next to the Steer. Above, equipment and cooked patties at Beyond Meat.

The Problem with Fake Meat

It might be possible to create a burger that helps the environment and improves your health. But will it taste good enough to win over the masses?

By Corby Kummer

People want burgers. It seems hardwired. You can read Richard Wrangham's *Catching Fire* to learn how man evolved into a thinking primate by learning to cook the animals he killed. You can talk to the stylish proprietor of a leading cooking school in Japan, who co-owns an artful Manhattan sushi restaurant. What does he find the most efficient fuel for his triathlon training? A couple of McDonald's quarter-pounders a day.

Vegetarian and vegans want burgers. Walter Robb, co-CEO of Whole Foods, says that from the time he started a health-food store in the Northern California of the late 1970s, he had to sell tofu, seitan, and anything else that could be made to look like meat but wasn't. "The stuff sells," he says simply. Entire books are dedicated to veggie burgers, even if they all taste like overseasoned, underhydrated corrugated cardboard.

Of course, there are rational reasons not to eat meat. You can probably recite them along with Ethan Brown, a strapping 6-foot-5 vegan who sold his house in Washington, D.C., and raided his family's savings accounts to fund a startup called Beyond Meat. Because raising livestock is such an inefficient use of land and water, he thought that making soy "chicken" strips and vegetable-protein "Beast" patties would be an even better way to improve the environment than creating fuel cells, the career he abandoned. Along the way he signed up Bill Gates and Twitter founders Biz Stone and Evan Williams as investors. It's hard, in fact, to find a tech billionaire who hasn't invested in a protein alternative that aims to stamp out factory farming. They all recognize the realities of the market: everybody buys burgers. "Meat is such a macho thing," Williams says.

I eat meat. It's hard to be a restaurant reviewer, as I am, without eating meat. I like to think I'm less culpable if it has been raised with care, killed humanely (not something whose meaning I'm so clear on, though from the time I started writing about food I've watched chickens, lambs, and cattle be killed and butchered in farms and factories), and sold at a price that allows fair wages to everyone involved with its production. But I have never tried to delude myself that more than the tiniest fraction of people who want meat can afford to keep these illusions of enlightenment alive.

The problem is that the new alternatives are—in the words of one tech billionaire who isn't sold on the idea, Nathan Myhrvold (a trained chef and author of the encyclopedic *Modernist Cuisine*)—essentially "slightly better Tofurkey." So why bother? This was the question on my mind when I headed to Beyond Meat's office in El Segundo, California. Why kill yourself to produce a not-quite-rubber burger? Why not just make something new?

Beast Burger, $5.89
Beastly Sliders, $5.89
Beyond Beef, $4.61
Beyond Chicken, $4.55
(average retail prices)

Cut Steakhouse
Beverly Hills, California

The Steer

Beyond Meat can count itself a tech company, in that it began when Brown plowed through scientific papers to find the university researchers who were doing the work likeliest to advance the T in TVP—textured vegetable protein, which has generally had a consistency somewhere between modeling clay and latex sponge. Texture, Brown thought, was the key to a better meat substitute. He also wanted to vary the V: most TVP means soy, in a world where many people want to avoid genetically modified organisms and almost all soy is a GMO. His assumption was that the flavor challenge had been cracked by chemists working from the late 1960s through the '80s—a golden era for experimentation in processed food, when instruments to measure flavor were being invented and refined, multinational flavoring companies were racing to develop new molecules, and cranks hadn't started talking about eating only what your grandmother ate.

The El Segundo offices, on a street in a quiet beachside neighborhood with strip malls, seem less like a tech startup than like the laid-back domain of amiable tinkerers. The essential research machine is a clunky-looking extruder Brown calls "the Steer," to point out its efficiency at converting feed to "meat." Product developers scoop white soy and pea protein meal that looks like animal feed from white plastic buckets into one end of the machine, along with water, and grab strips out of the extruder to see how moist and tearable they feel. A young man at a walnut-Formica table uses an eyedropper to squeeze rust-colored liquid, colored with turmeric, onto a Beast patty to see if he can get the liquid to stay in the burger during cooking and create the look of myoglobin, a protein in muscle cells; the paper plates under the cooked patties stubbornly keep getting stained. The test kitchen has an open pantry of wire shelves filled with spices and peppers and starch powders only slightly more abstruse than what you'd find in a supermarket, and a stove that could have come out of a rental apartment.

Dave Anderson, a friendly, slightly shaggy chef, ran a popular vegan restaurant in Los Angeles, where he was particularly proud of his multistep portobello mushroom bacon and his seitan poached in mushroom broth ("You

could tear it like filet mignon"). Brown may have assumed that flavor was the easy part and texture the hard part, but Anderson has learned from trial and error that both are still high hills to climb. Meat, says Don Mottram, an emeritus professor of food chemistry at the University of Reading, is the hardest problem for flavor-chemical companies to solve. Because of its complex structure, he says, meat develops flavor at different rates as fat, muscle, and bone successively cook. Mottram spent decades investigating meat flavor and in particular the Maillard reaction, the caramelization of carbohydrates that releases hundreds or thousands of compounds during cooking.

Because of its complex structure, meat is the hardest thing to imitate with flavorings and textured vegetable protein.

Anderson approaches flavor like the cook he is: by constantly experimenting with the proportions of ingredients. He gamely warms up some Beyond Chicken "lightly seasoned" strips, Beyond Beef "beefy crumbles," and a Beast Burger for me to taste against their real-meat counterparts—something that he and the rest of the flavor developers, including the diehard vegans, regularly do (they figure that giving fellow vegans better alternatives will make up any lost karma).

I'm impressed by the "lightly seasoned" strips, insofar as they bear a strong similarity to my Hungarian grandmother's Saturday lunch of reboiled chicken from her Friday-night chicken soup. She used garlic and onion in everything, too much salt, and usually some dried or fresh parsley. So does Anderson. What took her hours of simmering for a particular waterlogged yet dry, chewy texture, Beyond Meat achieves by tossing pieces of extruded soy protein into a flavored brine under a vacuum, so the liquid and flavor will penetrate better. The chicken strips out of a bag from Tyson or a similar mass-market supermarket brand—the standard Anderson says he is aiming for—do have a meatier flavor. But only slightly. Their chewy-fibrous texture is more unmistakably meaty than that of the Beyond Meat strips, though the strips are pretty close and getting closer. Tim Geistlinger, who's in charge of R&D, lets me sample a new batch of "chicken" strips extruded from the Steer, which have a more variegated and branching striation than the current version. With better hydration, the newly configured strips will be possible to confuse with something out of a Tyson bag. I went through the better part of a bag of Beyond Meat strips without really thinking about it. And I'd certainly rather eat what Beyond Meat extrudes than what Tyson packages.

The Beast is more problematic. It has to be tarted up with a lot of seasoning—more onion and garlic, paprika, mesquite, sugar—to cover the taste of the nutrient powder it contains so that Brown can claim in TED-talk tones that it has more iron and protein than the same amount of ground beef, more omega-3s than the same amount of salmon. What the Beast doesn't have is enough hydration to keep you from needing a good bit of liquid to get it down. When I saw Brown eat one, he added the emollients of ketchup, sliced tomato, and iceberg lettuce.

What's most striking is not how close these products are to supermarket chicken strips and ground beef but how debased our own flavor sense has

A Beast Burger's relative dryness can be addressed with typical burger fixings.

become. If Bill Gates and other luminary investors in Beyond Meat can be fooled, as they say they have been, it may be more because of what they're used to than what actual chicken or steak tastes like.

After my taste-through, I went to Cut, one of Los Angeles's most expensive steakhouses, in the Beverly Wilshire hotel. There's nothing like a steak with the intramuscular marbling fat that bastes every bite of a bone-in porterhouse: tender loin with just enough chew not to seem rancid, sinew and cartilage for texture, and a heavy fat cap that is like a food group of its own. Beyond Meat and its rivals are decades away from anything like that.

But as for the kobe sliders that came as a giveaway at Cut after my table ordered enough steak to make it worth the restaurant's while: once you scrape away the ashy char and ignore the house-made ketchup and freshly baked brioche bun, the chewy gristle isn't so far from the dry, flavor-free crumbs of supermarket ground beef Anderson plunked down beside the Beast. Plain ground beef is dismal. With some essential work on flavor and moisture, Anderson and Geistlinger will be able to get beyond the cooked-dog-food appearance of the Beast. They might even perfect the Salisbury steak, that staple of school cafeterias, that Anderson says he can imagine achieving in his lifetime (he doesn't mention the school-cafeteria part), or the skinless chicken breast that both men think might not be far down the road.

Another alternative—test-tube meat, also known as cultured meat, in vitro meat, and lab meat—is probably decades off, despite the introduction of a $332,000 burger at a London press conference in August 2013. The pinkish ground meat had been produced in a Maastricht University lab directed by Mark Post, a vascular biologist and surgeon: it consisted of billions of cells cultured from skeletal muscle cells taken from one beef neck, nourished in a warm broth of synthetic nutrients and cow-fetus serum. To get the cells to grow into myotubes, the building blocks of muscle fiber, the researchers reduce the serum in the broth, which causes the cells to stop dividing and fuse. Then they suspend the cells in a gel surrounding a central column that allows them to align and form muscle fibers. For the scaffold, Post and others first used Velcro and then searched out biodegradable options. At the live-streamed tasting, the testers reported that the burger tasted almost like a real one, but not as juicy and "surprisingly crunchy." (The burger backer was Sergey Brin.)

Somewhat more practical-minded researchers based in Brooklyn, New York, are aiming to produce cultured meat at a company called Modern Meadow (the names of these companies, you will have noticed, border on the Orwellian). Gabor Forgacs, a theoretical physicist who changed midcareer to developmental biology, and his son, Andras, are incubating beef cells and mixing them with pectin and spices to create a range of products, including "baked steak chips." Their original company, Organovo, intended to produce living tissue for drug testing; food seemed to be an equally achievable goal. Of course, Modern Meadow has its own Silicon Valley angel: Peter Thiel.

In theory, cultured meat can be scaled and may offer something closer to real meat than any other inventions in the works. By its nature, it would offer the complex flavors of meat. But it is still in the basic-research phase. The problems are many: scientists must figure out how to build intramuscular fat, sinew, cartilage, and even bone, and a structure to mimic veins and blood vessels that will keep the cells fed so they don't become gangrenous. The work is so expensive that the steps forward are likely to come from trying to produce organs for transplant—which are "worth millions of dollars a pound instead of $10 a pound," as Myhrvold points out.

Truly new

None of this will do much for people who care about cuisine. Fooling more people by coming closer to debased industrial meat will hardly elevate America's palate.

Admittedly, none of these companies is aiming to do that: people on the frontiers of flavor are not their intended audience. But for people who do want to be on the frontiers, some of the new research could result in actual improvements.

What I'm interested in seeing is how cooks will use these companies' protein-isolation techniques to create entirely new textures. Two ethereal dishes pointed the way for me. I tried them during a competition among practitioners of washoku, a Japanese cooking philosophy that glorifies umami with results from the simple to the exquisite. One was a pyramid of trembling, subtle sesame tofu, a Kyoto specialty of Buddhist monks. Flanlike, with the musky flavor of toasted sesame and light soy, it didn't attempt to be anything but delicate, and it was unlike any tofu I'd had anywhere, including Koreatown restaurants that make fresh batches every few hours.

The other was a small white bowl of luminescent white tofu as reconceived by Rene Redzepi, Lars Williams, and their staff when they set up an outpost of the Danish restaurant Noma in Tokyo last January and February. It was the one classic Japanese dish they dared try to make, Williams explained one night as we stood watching the kitchen crew. Tiny corkscrews of soft, beige grated unripe walnut coated it like snow; an emerald-green herb sauce lay at the bottom. The tiny cube of tofu didn't quite taste of milk or soy, though it was reminiscent of both; it was silken air, the clear expression of a passing if intense notion of what fresh tofu could be. In the hands of cooks capable of that kind of imagination and high-wire skill, pea-protein isolates—even fortified with omega-3s and iron—can be the way to save the world and keep it safe for culinary invention, too.

Corby Kummer is an editor at the Atlantic *and reviews restaurants for* Boston *and* Atlanta *magazines.*

Survival in the Age of Spotify

Two rock musicians find flaws—and hope—in a book that suggests how artists can earn a decent living even after free online access to music has ravaged the business.

By Aimee Mann and Ted Leo

Mann: Ted, we are both intimately affected by the issues discussed in Cory Doctorow's book *Information Doesn't Want to Be Free*. I was very apprehensive about how to approach it. I thought, "If I'm going to be reading a bunch of suggestions about how I can tweet for couches to sleep on after shows, I'm going to be really depressed." And in fact, in the beginning of the book there's a lot of that language we're familiar with that comes across as: "Those artists out there who were doing okay by the old systems and now are flailing—too bad! Sorry, lamplighters! Too bad you couldn't keep up, buggy-whip makers!"

Leo: I've been frustrated that artists are not allowed to have a nuanced position. We seem to get shoehorned into the extreme camps of "completely pro–free culture" or "completely anti-technology." When the record labels or publishing houses get together to hammer out deals with intermediaries (be they ISPs, streaming services, or digital storefronts), we are rarely, if ever, invited to the table and have no choice but to react to what's already been decided. As someone who loves the Internet but hopes for more creative solutions to "free vs. paid" than the binary extremes we seem to be forced to choose between, I am pretty tired of that language. It's off-putting to have a book that purports to have your best interests at heart open with so much of that language of conflict.

Mann: Right, and I'm not opposed to the usual solutions along the lines of "Tour more and sell T-shirts," because that's what we've been doing for decades. Nor am I opposed to the idea of musicians raising money through crowdfunding. I like all the creative energy you can put into premiums, videos, etc. You have to find a way to creatively self-market. But the need to keep turning up the volume and increasing the spectacle, clamoring for attention online, sometimes starts to sink to the level of an old-timey PR stunt, along the lines of flagpole sitting, and that isn't appropriate for certain kinds of art or artists.

Leo: Indeed. And it isn't guaranteed to be effective. I am truly moved seeing artists create mutually beneficial exchanges with their fans through crowdsourcing. But it's worth considering, also, that in this particular model, fans' investment in the person and personality of the artist drives their loyalty almost as much as the music. It's an awesome thing to witness, but I'm not sure a market based on that model could even sustain more than a handful of successes. So it's important to be careful of saying that this is the future—the only way to sustainability.

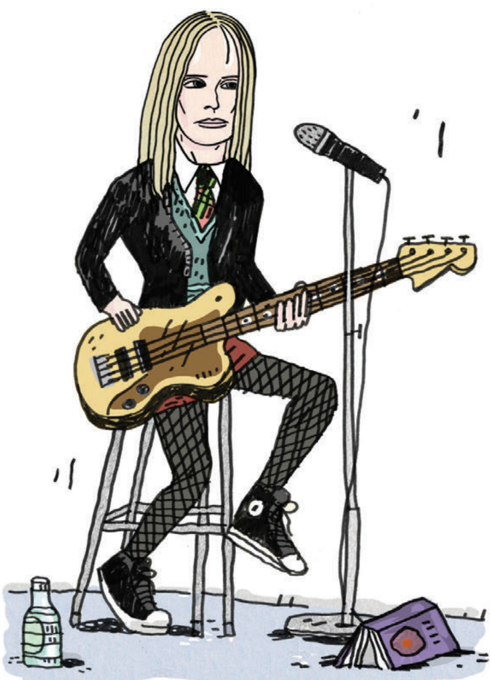

MARK TODD

Mann: There's no one-size-fits-all solution. Even identifying the problems facing artists as universal is wrong. Circumstances change, and artists are usually some of the most adaptable people on the planet.

Leo: The fact that we both felt put on the defensive by the tone early on is a flaw in the book itself. Because as Doctorow delves into the history of copyright law and digital locks later on, I found myself rapt and in pretty close alignment with his assessment of how negatively those locks affect us broadly—culturally, societally.

Mann: Oh, he hates digital locks. Do not bring a digital lock around Cory Doctorow!

Leo: And I agree with him. I hated Apple's digital rights management from day one. I remember trying to DJ my sister's wedding and wanting to play a few things from my library off someone else's iTunes, having to deauthorize and reauthorize computers, and manage iPods and hard drives, etc.—First World problems, maybe, but frustrating impediments to just enjoying music, nonetheless. And probably more importantly, Doctorow opened my eyes to some of the more secretive and more nefarious ways digital locks are operating.

Mann: Yes. One of the most frightening things I learned is how locks are baked into our machines. We're actually ceding control and choice and allowing our computers to do things we don't even know.

At this point I have to go back to your point about the initial tone being a problem with the book, because it reads a bit like notes from a lecture or a TED talk. It's more a collection of essays with a lot of asides and sidebars, but the through line is hard to discern.

Leo: There are passages that, when I think about them being spoken to a live audience, sound better than they read on a page.

Mann: It makes me wonder if the dynamic in publishing is similar to the one in the music world that makes people say "Tour more"—one that's placing authors in the position of finding a better (or at least supplementary) business model in generating speaking fees than in selling books.

To jump ahead, then, to what I'd call the third section, in which he finally lays out his idea for getting people paid for the content they create, I was interested and pleasantly surprised at how artist-friendly it seems to be.

Leo: I agree. For the unfamiliar, in general, the performers' rights organizations (ASCAP, BMI, SESAC) administer the copyrights for a vast collection of musical work and license the use of these catalogues.

Mann: The license takes a percentage of the gross revenue of, for example, an FM radio station, and out of that sum disburses royalties to composers.

Leo: And in principle, it works similarly with the streaming services, but it's very convoluted. Pandora, for example, operates under licenses that pay both performers and composers, though in a proportion that, unlike radio licenses, leaves composers feeling shorted. To make it even more confusing, Pandora also does direct deals with rights holders (usually labels, but theoretically composers as well). Spotify and YouTube, however, have their own rates that they've negotiated directly with the big three record labels (which also are investors in Spotify), often leaving independent artists and labels in a "take it or leave it" situation that affects their bottom line more severely. At the same time, it sets a high standard for entry into the streaming market that many innovators can't afford.

Information Doesn't Want to Be Free
By Cory Doctorow
McSweeney's, 2014

Doctorow's idea is essentially that the sorts of blanket licenses that are used in radio could be fitted to Internet streaming or downloading, with artists actually guaranteed a (rather large) percentage of the pay-out pie. Innovation would be easier, with services—especially startups—able to focus on delivering ever better experiences instead of having to spend inordinate amounts of time and money negotiating with big labels and publishing houses.

It hasn't gotten a lot of play, but the idea's been around for a while. Jaron Lanier has written about his vision for having nano-payments sent to every originator of content that is referenced or used anywhere. Doctorow's idea almost seems like the version that the big labels, publishing houses, and intermediaries would eventually water Lanier's plan down to!

Mann: Ha! And in fact, just this winter the U.S. Copyright Office issued some new recommendations on exactly this—how music licensing can be updated to better reflect the new realities of how we create and experience music. I do like it as a creative solution that might actually encourage art and innovation for people just jumping in or even decades into a career.

Leo: Yes. Look—as dumb as this may sound, I am an old punk rocker and fan of futurism and still find a lot to get excited about in the Silicon Valley mantra of "Break shit, apologize later." I understand the frustrations with "The System," writ large, and digital locks, specifically, and the creative spirit that drives people to want to "break" them, but ultimately, it's because we're makers. I appreciate where Doctorow eventually lands in this book, and I think there's something to be said for the hope of being able to focus more on actually "making shit" when we no longer need to "break shit."

Aimee Mann and Ted Leo play in a duo called the Both.

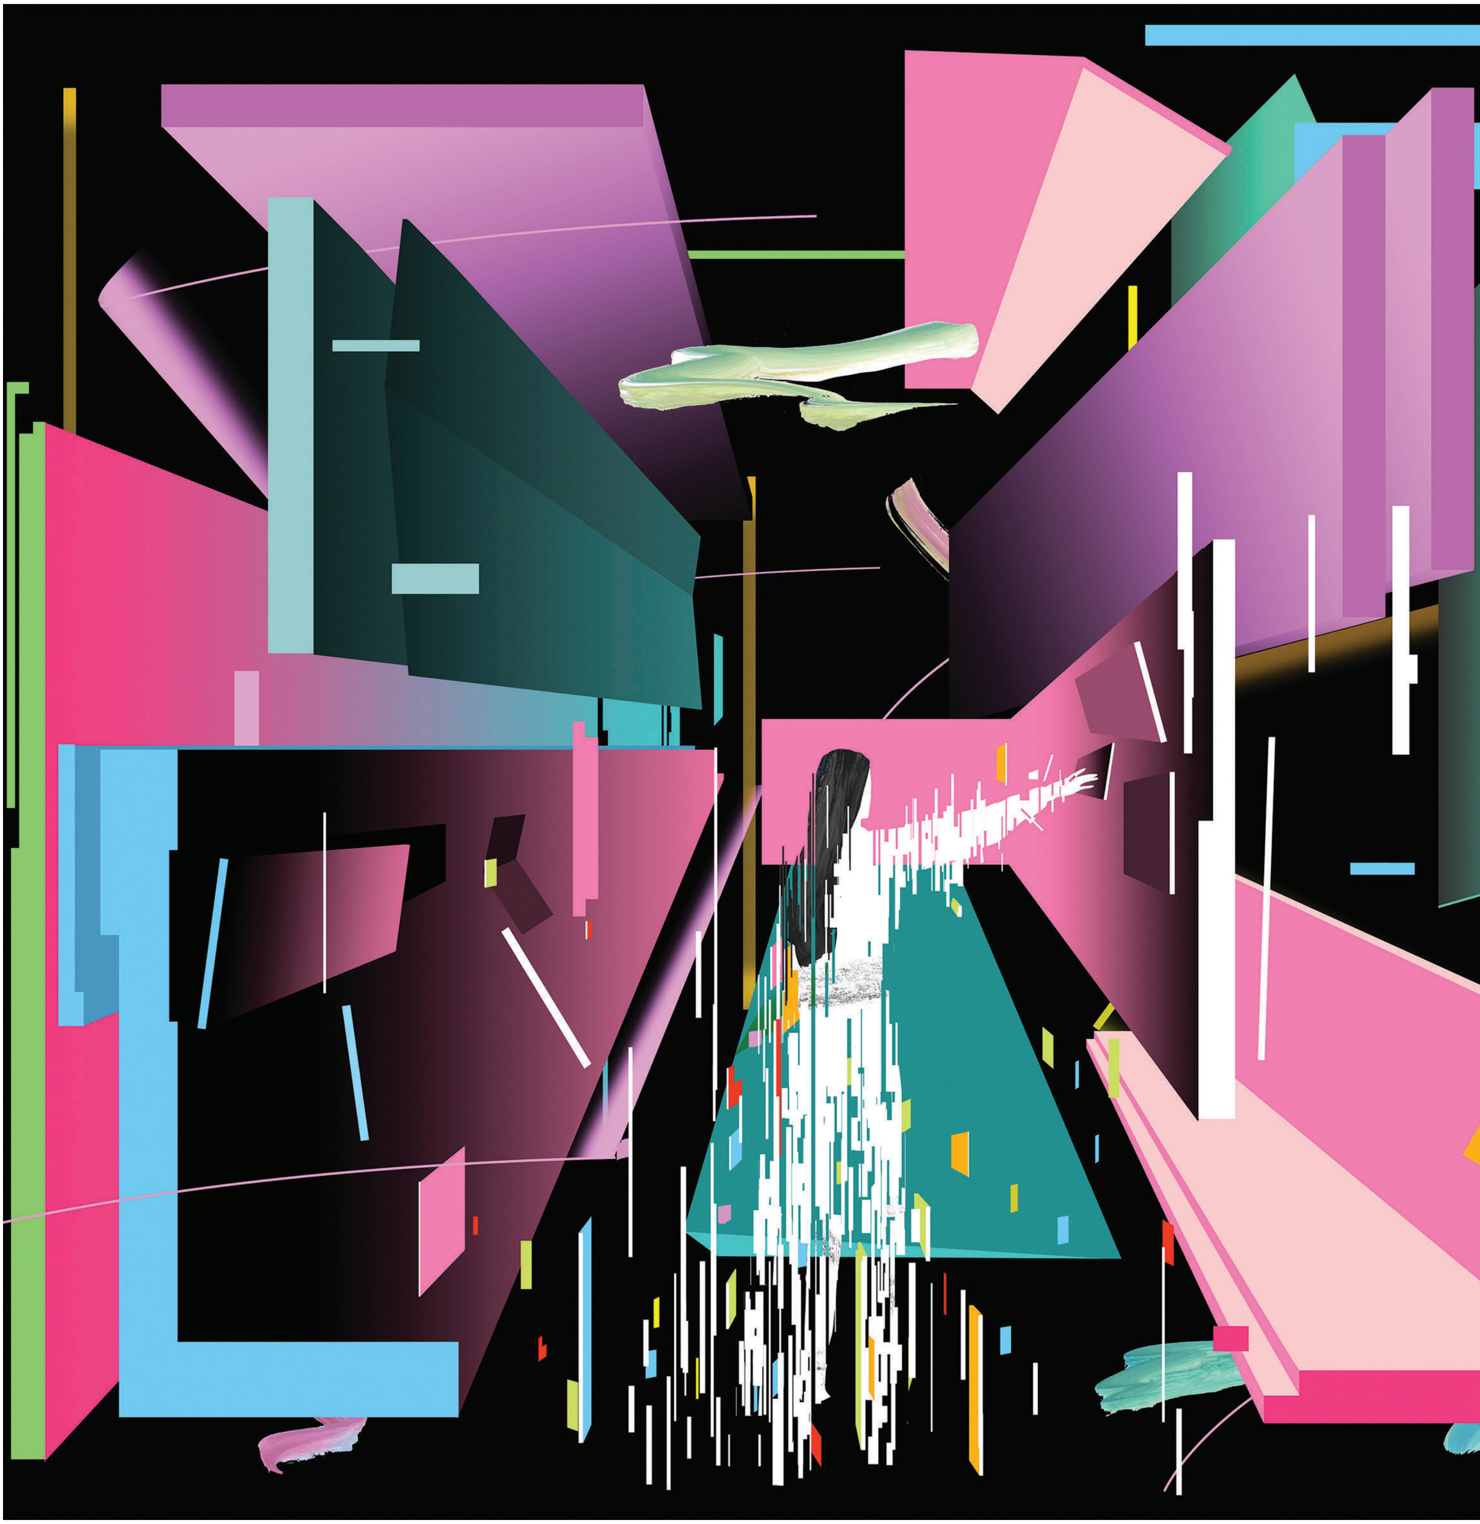

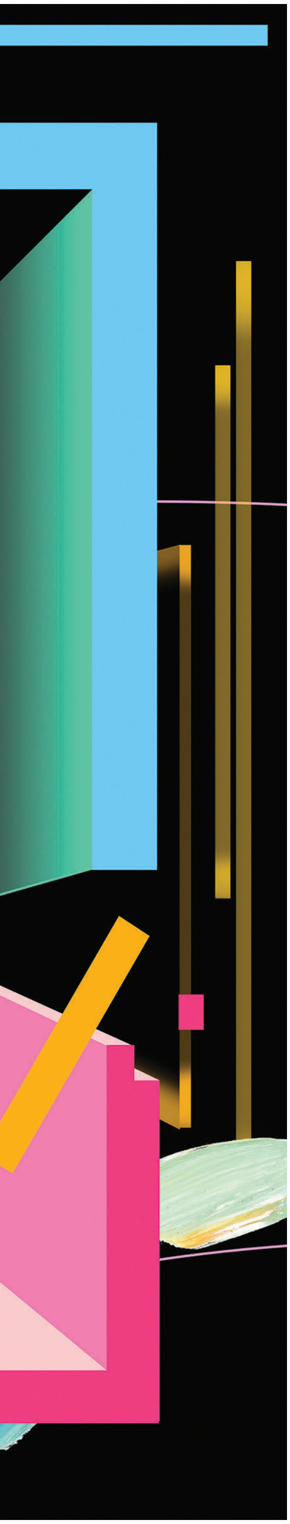

JON HAN

Toolkits for the Mind

Programming languages shape the way their users think—which helps explain how tech startups work and why they are able to reinvent themselves.

By James Somers

Babel-17
By Samuel R. Delany
1966

Real World OCaml
By Yaron Minsky et al.
O'Reilly, 2013

PHP

Hack

Scala

When the Japanese computer scientist Yukihiro Matsumoto decided to create Ruby, a programming language that has helped build Twitter, Hulu, and much of the modern Web, he was chasing an idea from a 1966 science fiction novel called *Babel-17* by Samuel R. Delany. At the book's heart is an invented language of the same name that upgrades the minds of all those who speak it. "Babel-17 is such an exact analytical language, it almost assures you technical mastery of any situation you look at," the protagonist says at one point. With Ruby, Matsumoto wanted the same thing: to reprogram and improve the way programmers think.

It sounds grandiose, but Matsumoto's isn't a fringe view. Software developers as a species tend to be convinced that programming languages have a grip on the mind strong enough to change the way you approach problems—even to change which problems you think to solve. It's how they size up companies, products, their peers: "What language do you use?"

That can help outsiders understand the software companies that have become so powerful and valuable, and the products and services that infuse our lives. A decision that seems like the most inside kind of inside baseball—whether someone builds a new thing using, say, Ruby or PHP or C—can suddenly affect us all. If you want to know why Facebook looks and works the way it does and what kinds of things it can do for and to us next, you need to know something about PHP, the programming language Mark Zuckerberg built it with.

Among programmers, PHP is perhaps the least respected of all programming languages. A now canonical blog post on its flaws described it as "a fractal of bad design," and those who willingly use it are seen as amateurs. "There's this myth of the brilliant engineering that went into Facebook," says Jeff Atwood, co-creator of the popular programming question-and-answer site Stack Overflow. "But they were building PHP code in Windows XP. They were hackers in almost the derogatory sense of the word." In the space of 10 minutes, Atwood called PHP "a shambling monster," "a pandemic," and a haunted house whose residents have come to love the ghosts.

Most successful programming languages have an overall philosophy or set of guiding principles that organize their vocabulary and grammar—the set of possible instructions they make available to the programmer—into a logical whole. PHP doesn't. Its creator, Rasmus Lerdorf, freely admits he just cobbled it together. "I don't know how to stop it," he said in a 2003 interview. "I have absolutely no idea how to write a programming language—I just kept adding the next logical step along the way."

Programmers' favorite example is a PHP function called "mysql_escape_string," which rids a query of malicious input before sending it off to a database. (For an example of a malicious input, think of a form on a website that asks for your e-mail address; a hacker can enter code in that slot to force the site to cough up passwords.) When a bug was discov-

ered in the function, a new version was added, called "mysql_*real*_escape_string," but the original was not replaced. The result is a bit like having two similar-looking buttons right next to each other in an airline cockpit: one that puts the landing gear down and one that puts it down *safely*. It's not just an affront to common sense—it's a recipe for disaster.

Yet despite the widespread contempt for PHP, much of the Web was built on its back. PHP powers 39 percent of all domains, by one estimate. Facebook, Wikipedia, and the leading publishing platform WordPress are all PHP projects. That's because PHP, for all its flaws, is perfect for getting started. The name originally stood for "personal home page." It made it easy to add dynamic content like the date or a user's name to static HTML pages. PHP allowed the leap from tinkering with a web*site* to writing a Web *application* to be so small as to be imperceptible. You didn't need to be a pro.

PHP's get-going-ness was crucial to the success of Wikipedia, says Ori Livneh, a principal software engineer at the Wikimedia Foundation, which operates the project. "I've always loathed PHP," he tells me. The project suffers from large-scale design flaws as a result of its reliance on the language. (They are partly why the foundation didn't make Wikipedia pages available in a version adapted for mobile devices until 2008, and why the site didn't get a user-friendly editing interface until 2013.) But PHP allowed people who weren't—or were barely—software engineers to contribute new features. It's how Wikipedia entries came to display hieroglyphics on Egyptology pages, for instance, and handle sheet music.

You wouldn't have built Google in PHP, because Google, to become Google, needed to do exactly one thing very well—it needed search to be spare and fast and meticulously well engineered. It was made with more refined and powerful languages, such as Java and C++. Facebook, by contrast, is a bazaar of small experiments, a smorgasbord of buttons, feeds, and gizmos trying to capture your attention. PHP is made for *making*—for cooking up features quickly.

You can almost imagine Zuckerberg in his Harvard dorm room on the fateful day that Facebook was born, doing the least he could to get his site online. The Web moves so fast, and users are so fickle, that the only way you'll ever be able to capture the moment is by being first. It didn't matter if he made a big ball of mud, or a plate of spaghetti, or a horrible hose cabinet (to borrow from programmers' rich lexicon for describing messy code). He got the thing done. People could use it. He wasn't thinking about beautiful code; he was thinking about his friends logging in to "Thefacebook" to look at pictures of girls they knew.

The programming language PHP created and sustains Facebook's move-fast, hacker-oriented corporate culture.

Today Facebook is worth more than $200 billion and there are signs all over the walls at its offices: "Done is better than perfect"; "Move fast and break things." These bold messages are supposed to keep employees in tune with the company's "hacker" culture. But these are precisely *PHP's* values. Moving fast and breaking things is in fact so much the essence of PHP that anyone who "speaks" the language indelibly thinks that way. You might say that the language itself created and sustains Facebook's culture.

The secret weapon

If you wanted to find the exact opposite of PHP, a kind of natural experiment to show you what the other extreme looked like, you couldn't do much better than the self-serious Lower Manhattan headquarters of the financial trading firm Jane Street Capital. The 400-person company claims to be responsible for roughly 2 percent of daily equity trading volume in the United States.

When I meet Yaron Minsky, Jane Street's head of technology, he's sitting at a desk with a working Enigma machine beside him, one of only a few dozen of the World War II code devices left in the world. I would think it the clear winner of the contest for Coolest Secret Weapon in the Room if it weren't for the way he keeps talking about an obscure programming language called OCaml. Minsky, a computer science PhD, convinced his employer 10 years ago to rewrite the company's entire trading system in OCaml. Before that, almost nobody used the language for actual work; it was developed at a French research institute by academics trying to improve a computer system that automatically proves mathematical theorems. But Minsky thought OCaml, which he had gotten to know in grad school, could replace the complex Excel spreadsheets that powered Jane Street's trading systems.

OCaml's big selling point is its "type system," which is something like Microsoft Word's grammar checker, except that instead of just putting a squiggly green line underneath code it thinks is wrong, it won't let you run it. Programs written with a type system tend to be far more reliable than those written without one—useful when a program might trade $30 billion on a big day.

Minsky says that by catching bugs, OCaml's type system allows Jane Street's coders to focus on loftier problems. One wonders if they have internalized the system's constant nagging over time, so that OCaml has become a kind of Newspeak

that makes it impossible to think bad thoughts.

The catch is that to get the full benefits of the type checker, the programmers have to add complex annotations to their code. It's as if Word's grammar checker required you to diagram all your sentences. Writing code with type constraints can be a nuisance, even demoralizing. To make it worse, OCaml, more than most other programming languages, traffics in a kind of deep abstract math far beyond most coders. The language's rigor is like catnip to some people, though, giving Jane Street an unusual advantage in the tight hiring market for programmers. Software developers mostly join Facebook and Wikipedia in spite of PHP. Minsky says that OCaml—along with his book *Real World OCaml*—helps lure a steady supply of high-quality candidates. The attraction isn't just the language but the kind of people who use it. Jane Street is a company where they play four-person chess in the break room. The culture of competitive intelligence and the use of a fancy programming language seem to go hand in hand.

Google appears to be trying to pull off a similar trick with Go, a high-performance programming language it developed. Intended to make the workings of the Web more elegant and efficient, it's good for developing the kind of high-stakes software needed to run the collections of servers behind large Web services. It also acts as something like a dog whistle to coders interested in the new and the difficult.

Growing up

In late 2010, Facebook was having a crisis. PHP was not built for performance, but it was being asked to perform. The site was growing so fast it seemed that if something didn't change fairly drastically, it would start falling over.

Switching languages altogether wasn't an option. Facebook had millions of lines of PHP code, thousands of engineers expert in writing it, and more than half a billion users. Instead, a small team of senior engineers was assigned to a special project to invent a way for Facebook to keep functioning without giving up on its hacky mother tongue.

One part of the solution was to create a piece of software—a compiler—that would translate Facebook's PHP code into much faster C++ code. The other was a feat of computer linguistic engineering that let Facebook's programmers keep their PHP-ian culture but write more reliable code.

The rescue squad did it by inventing a dialect of PHP called Hack. Hack is PHP with an *optional* type system; that is, you can write plain old quick and dirty PHP—or, if you so choose, you can tie yourself to the mast, adding annotations to let the type system check the correctness of your code. That this type checker is written entirely in OCaml is no coincidence. Facebook wanted its coders to keep moving fast in the comfort of their native tongue, but it didn't want them to have to break things as they did it. (Last year Zuckerberg announced a new engineering slogan: "Move fast with stable infra," using the hacker shorthand for the infrastructure that keeps the site running.)

Around the same time, Twitter underwent a similar transformation. The service was originally built with Ruby on Rails—a popular Web programming framework created using Matsumoto's Ruby and inspired in large part by PHP. Then came the deluge of users. When someone with hundreds of thousands of followers tweeted, hundreds of thousands of other people's timelines had to be immediately updated. Big tweets like that would frequently overwhelm the system and force engineers to take the site down to allow it to catch up. They did it so often that the "fail whale" on the company's maintenance page became famous in its own right. Twitter stopped the bleeding by replacing large pieces of the service's plumbing with a language called Scala. It should not be surprising that Scala, like OCaml, was developed by academics, has a powerful type system, and prizes correctness and performance even at the expense of the individual programmers' freedom and delight in their craft.

Much as startups "mature" by finally figuring out where their revenue will come from, they can cleverly use the power of programming languages to manipulate their organizational psychology. Programming-language designer Guido van Rossum, who spent seven years at Google and now works at Dropbox, says that once a software company gets to be a certain size, the only way to stave off chaos is to use a language that requires more from the programmer up front. "It feels like it's slowing you down, because you have to say everything three times," van Rossum says. That is why many startups wait as long as they can before making the switch. You lose some of the swaggering hackers who got you started, and the possibility that small teams can rush out new features. But a more exacting language helps people across the company understand one another's code and gives your product the stability needed to be part of the furniture of daily life.

That software startups can perform such maneuvers might even help explain why they can be so powerful. The expanding reach of computers is part of it. But these companies also have a unique ability to remake themselves. As they change and grow, they can do more than just redraw the org chart. Because they are built in code, they can do something far more drastic. They can rewire themselves, their culture, the very way they think.

James Somers is a writer and programmer in New York. He works at Genius.com.

Demo

Polina Anikeeva

A Swiss Army Knife for Neuroscience

Neural probes that combine optics, electronics, and drugs could help unlock the secrets of the brain.

By Antonio Regalado
Photographs by Ken Richardson

Various powerful new tools for exploring and manipulating the brain have been developed over the last few years. Some use electronics, while others use light or chemicals.

At one MIT lab, materials scientist Polina Anikeeva has hit on a way to manufacture what amounts to a brain-science Swiss Army knife. The neural probes she builds carry light while collecting and transmitting electricity, and they also have tiny channels through which to pump drugs.

That's an advance over metal wires or silicon electrodes conventionally used to study neurons. Anikeeva makes the probes by assembling polymers and metals into large-scale blocks, or preforms, and then stretching them into flexible, ultrathin fibers.

Multifunctional fibers offer new ways to study animal behavior, since they can record from neurons as well as

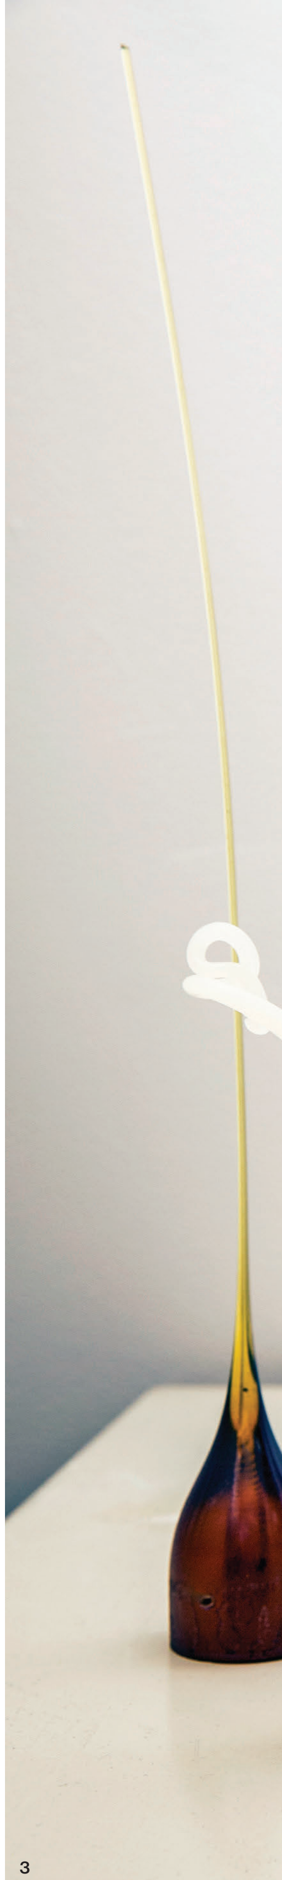

1 Blocks of polymers are the starting point for making a multifunctional neural probe. In a machine shop, patterns of conductive metal rods, transparent plastics, or hollow spaces will be added, creating a "preform."

2 The preform is loaded into this 12-foot-tall fiber-drawing tower.

3 A collection of preform leftovers after drawing. Indium-tin rods are visible in what's left of the preform at center.

4 Fiber is pulled from the furnace after being heated to 350 °C. A micrometer (red light) monitors the fiber's size.

5 Each preform is drawn into as much as one kilometer of fiber. It is now about 1/100th as thick as it was originally.

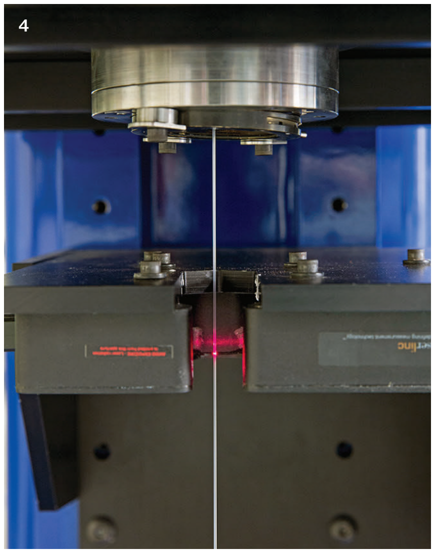

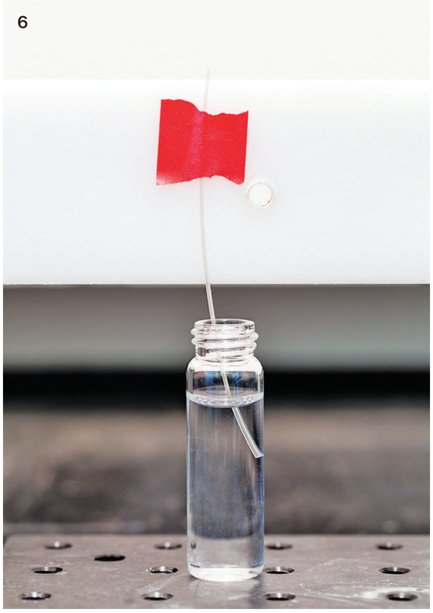

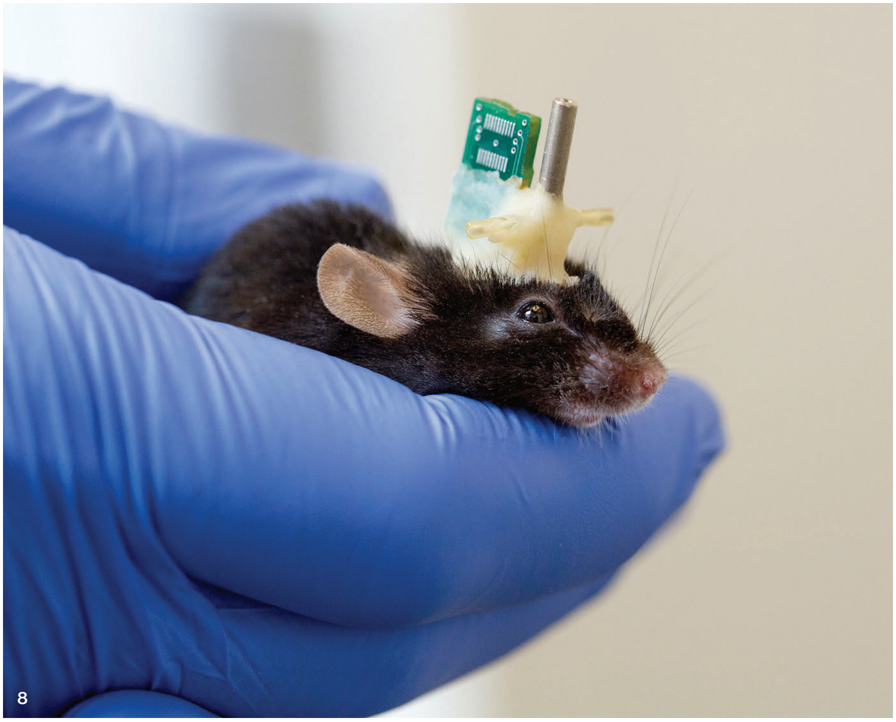

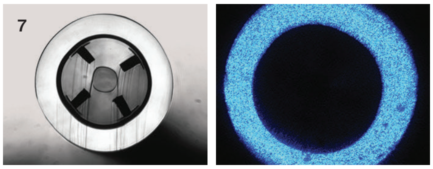

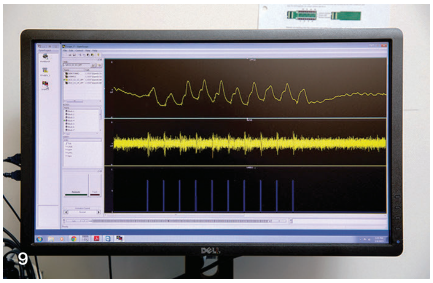

6 A fiber soaks in THF, a solvent, to remove a protective cladding.

7 A cross section of a 0.35-millimeter-wide fiber containing four electrodes, a fluid channel, and a ring-shaped waveguide. At right, light shines through the waveguide.

8 This mouse has a fiber implanted in its brain. Visible on its head are a circuit board, a port to introduce light, and two more to inject drugs.

9 Optically stimulating the mouse's brain produces the electrical activity recorded here.

stimulating them. New types of medical technology could also result. Imagine, as Anikeeva does, bionic wiring that bridges a spinal-cord injury, collecting electrical signals from the brain and transmitting them to the muscles of a paralyzed hand.

Anikeeva made her first multifunctional probe while studying at Stanford. It was crude: she simply wrapped metal wires around a glass filament. But this made it possible to combine standard electrode measurements with a new technology, optogenetics, in which light is fired at neurons to activate them or shut them down.

Now Anikeeva, a professor of materials science and engineering, makes probes using a fiber-drawing technology developed by another MIT researcher, Yoel Fink. It's based on the way silica is heated and pulled to form telecommunications fiber. But it works at lower temperatures, at which many useful polymers become soft enough to stretch.

Polymer fibers have a couple of important advantages. One is that they are flexible and mimic the physical properties of tissue. That could allow them to work longer than the stiff metal electrodes neuroscientists have relied on, permitting long-term studies in animals. The second feature of the fibers is that they can combine many functions. Probes made so far have incorporated as many as 36 microwires, optical waveguides, and hollow channels for carrying medicine. There's no reason not to incorporate sensors to measure temperature or pressure as well. Inside the body, the right materials and structures might even entice nerves to attach to the fibers, the way bone fuses to a hip implant.

The fiber-drawing process shrinks large patterns into microscopic ones, preserving the details. But there are challenges. The tiny wires and tubes have to be stripped, splayed, and soldered by hand to connect them to components such as a recording device a mouse wears on its head. That's quite a nightmare, says Andres Canales, a graduate student, who hopes to resolve the problem.

Will polymer bio-wires be what ultimately cures paralysis—say, by ferrying nerve signals across an injured spinal cord? "I think it will be a version of this technology, a more sophisticated version," says Anikeeva. "At least we are going to pursue this route." ■

50 Years Ago

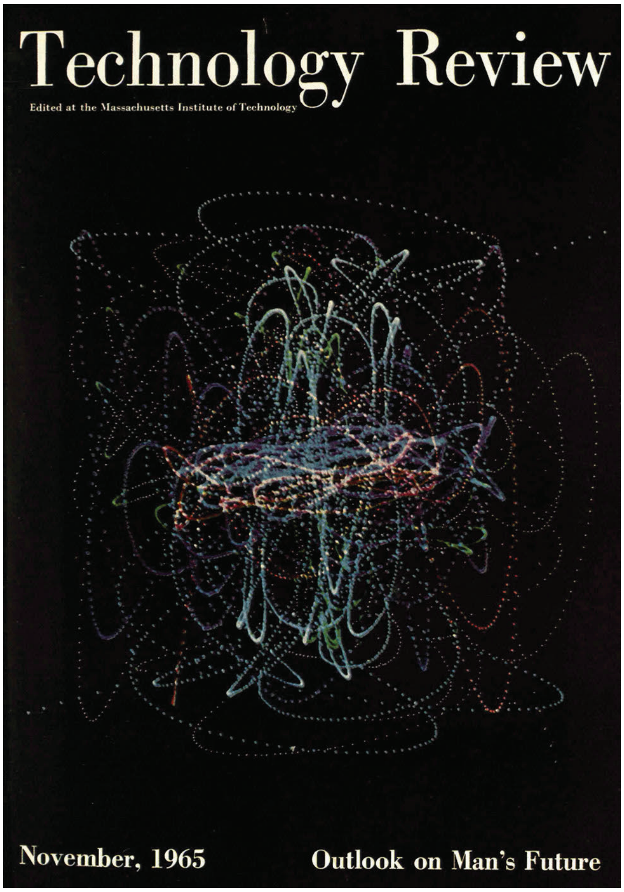

Eat Your Veggies

The long quest to replace meat.

"The present world population could be fed easily by the agricultural production of conventional foods if only there were time for crop yields in the less developed countries to catch up with population needs and if purchasing power kept pace. Since neither is occurring today, it is prudent to look to less conventional solutions to part of the food problem. There is no fixed nutritional requirement for the relatively costly sources of protein—milk, meat, and eggs. Legumes and the oilseed meals are acceptable alternatives. After the seeds of soy, cotton, sesame, peanut, and sunflower plants are commercially pressed, the resultant meal contains approximately 50 per cent protein, and the oil can be sold to pay much of the cost.

One-third of a properly processed oilseed meal mixed with two-thirds of cereal grain gives a mixture whose quality and concentration of protein is adequate for all human needs.

Algal proteins have attracted much attention, and the idea of a mechanical cow to extract edible protein from the juice of grass and leaves has been seriously advanced. The problem thus far has been poor palatability and relative high cost, but these are merely a challenge to food technologists.

Still more significant to those aware of the resources of modern science and technology is the steady progress toward new foods wholly synthetic in origin. The nutritional needs of the human body can be reduced to chemically known substances, and all of these can be synthesized or extracted from natural products. The needed vitamins and minerals are already so available and cheap as to constitute no problem. Of the eight essential amino acids needed instead of protein, the two for which there is commercial demand already are available by the ton, and there is no doubt that the others could be if demand were established.

At present it is prohibitively expensive to feed humans on diets of synthetic nutrients, and no suitable formulations are available, but highly palatable, attractive synthetic foods are within reach of modern technology."

Excerpted from "Food and Health: Contrasts Deepen," by Nevin S. Scrimshaw, head of MIT's Department of Nutrition and Food Science, in the November 1965 issue of Technology Review.

MIT Technology Review (ISSN 1099-274X), May/June 2015 Issue, Reg. U.S. Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Entire contents ©2015. The editors seek diverse views, and authors' opinions do not represent the official policies of their institutions or those of MIT. Periodicals postage paid at Boston, MA, and additional mailing offices. Postmaster: send address changes to MIT Technology Review, Subscriber Services, PO Box 5001, Big Sandy, TX 75755, or via the Internet at www.technologyreview.com/customerservice. Basic subscription rates: $41.94 per year within the United States; in all other countries, US$54. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051 Fort Erie, ON L2A 6C7. Printed in U.S.A. Audited by the Alliance for Audited Media

At one point, 40 percent of streetlights in Detroit didn't work. This made life even more difficult for a city that was already struggling.

The Public Lighting Authority of Detroit devised a plan to reilluminate the city. But finding a bank to finance the project during Detroit's bankruptcy was challenging. Citi stepped up and committed its own capital, which encouraged other investors. So far, thousands of new LED lights have been installed, lighting the way as a model for similar projects around the world.

For over 200 years, Citi's job has been to believe in people and help make their ideas a reality.

citi.com/progress

citi®
The World's Citi℠